

RESEARCH ARTICLE

INFLECT-DGNN: Influencer Prediction With Dynamic Graph Neural Networks

ELENA TIUKHOVA¹, EMILIANO PENALOZA², MARÍA ÓSKARSDÓTTIR³, BART BAESENS^{1,4},
MONIQUE SNOECK¹, (Member, IEEE), AND CRISTIÁN BRAVO²

¹Research Center for Information Systems Engineering (LIRIS), KU Leuven, 3000 Leuven, Belgium

²Department of Statistical and Actuarial Sciences, Western University, London, ON N6A 5B7, Canada

³Department of Computer Science, Reykjavik University, 102 Reykjavik, Iceland

⁴Department of Decision Analytics and Risk, University of Southampton, SO17 1BJ Southampton, U.K.

Corresponding author: Elena Tiukhova (elena.tiukhova@kuleuven.be)

This work was supported in part by ING Chair on Applying Deep Learning on Metadata as a Competitive Accelerator, in part by Canada Research Chairs Program under Grant CRC-2018-00082, in part by the Compute Ontario, in part by Calcul Québec, and in part by the Digital Research Alliance of Canada. The work of María Óskarsdóttir was supported by Icelandic Research Fund (IRF) under Grant 228511-051. The work of Cristián Bravo was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant Program under Grant RGPIN-2020-07114.

ABSTRACT Leveraging network information for predictive modeling has become widespread in many domains. Within the realm of referral and targeted marketing, influencer detection stands out as an area that could greatly benefit from the incorporation of dynamic network representation due to the continuous evolution of customer-brand relationships. In this paper, we present INFLECT-DGNN, a new method for profit-driven INFLuencer prEdiCTion with Dynamic Graph Neural Networks that innovatively combines Graph Neural Networks (GNNs) and Recurrent Neural Networks (RNNs) with weighted loss functions, synthetic minority oversampling adapted to graph data, and a carefully crafted rolling-window strategy. We introduce a novel profit-driven framework that supports decision-making based on model predictions. To test the framework, we use a unique corporate dataset with diverse networks, capturing the customer interactions across three cities with different socioeconomic and demographic characteristics. Our results show how using RNNs to encode temporal attributes alongside GNNs significantly improves predictive performance, while the profit-driven framework determines the optimal classification threshold for profit maximization. We compare the results of different models to demonstrate the importance of capturing network representation, temporal dependencies, and using a profit-driven evaluation. Our research has significant implications for the fields of referral and targeted marketing, expanding the technical use of deep graph learning within corporate environments.

INDEX TERMS Dynamic graph neural networks, graph isomorphism networks (GINs), graph attention networks (GATs), referral marketing, influencer prediction.

I. INTRODUCTION

Advances in analytics have enabled organizations to leverage data for effective decision-making and increased value creation [1]. A growing number of social interactions have become a valuable source for constructing networks that provide useful insights into interconnectivity patterns between individuals and can serve as a data source for real-world

The associate editor coordinating the review of this manuscript and approving it for publication was Xianzhi Wang¹.

applications such as fraud detection [2], recommender systems [3], fake news detection [4] and marketing [5]. Word of Mouth (WOM) has long been recognized as a powerful tool for customer influence in marketing to spread information effectively in customer networks [6]. Customers who are able to persuade others within their network can be considered influencers or opinion leaders [6]. Identifying and incentivizing them early helps to engage more customers in profitable initiatives by effectively spreading information across the network.

Typically, the technique used to incentivize customers who are considered influencers falls within the context of referral marketing, which leverages social connections to promote a product or service [7]. In this context, the topology of social networks is considered a major factor in the dissemination of information, with influencers being the best candidates to make referrals. To harness the power of social networks, it is essential to go beyond simply identifying opinion leaders. Anticipating the rise of potential influencers prior to their network enrollment, while also enabling re-influencing actions from past influencers who could repeatedly refer new customers, can serve as a significant competitive accelerator. Moreover, the value of predictive models depends heavily on the profit generated when the model's predictions are used in real marketing campaigns. Cost-sensitive decision making based on instance-dependent threshold tuning has proven to result in the highest savings [8]. This is achieved by complimenting a cost-insensitive classifier with cost-sensitive threshold optimization, which is known as the "predict-then-optimize" approach [8]. As a result, the use of influencer prediction models should not only be based on predictive performance, but also integrate a business view and evaluate the potential profit of the action, a property that is often overlooked in the existing studies. By doing so, businesses can effectively leverage the power of social networks to both expand their reach and drive financial growth, as well as more effectively target customers, reducing unnecessary spam and providing both a more nuanced approach to customer incentivizing and a better customer experience.

Network data is non-linear because it cannot be readily represented in Euclidean space. Therefore, the network topology should be extracted first to be used for downstream tasks. There are standard approaches for encoding network topology, such as neighborhood and centrality metrics, as well as collective inference algorithms [9]. With advances in computing power and deep learning, Graph Neural Networks (GNNs) have become a popular neural network architecture for learning from graph data [10]. GNNs have been successfully leveraged in different domains, including influencer detection [11], [12]. To account for dynamic networks, i.e. networks that change over time, Dynamic GNNs (DGNNs) came to the forefront of network analysis [13]. Öskarsdóttir et al. showed that incorporating the temporal component of the network improves predictive performance [14]. Therefore, GNNs can be used as the primary network representation learning technique and further complemented with the temporal component (e.g., recurrent neural networks (RNNs)) to account for the dynamic nature of social networks.

This paper presents INFLECT-DGNN, a novel method designed for influencer prediction in growing customer networks in a profit-driven manner. By leveraging the power of DGNNs to effectively learn from evolving network data, we solve a complex influencer prediction problem in dynamic networks with multiple relationships using

edge coloring, i.e., attributing specific characteristics to edges. Our method integrates two types of GNNs (Graph Attention Network (GAT) and Graph Isomorphism Network (GIN)) and Recurrent Neural Networks (RNNs). To learn from imbalanced data, INFLECT incorporates weighted loss functions and the Synthetic Minority Oversampling TEchnique (SMOTE) adapted for graphs [15]. INFLECT-DGNN also introduces a new profit-driven formula to enhance its practical applicability for decision making. Hence, we make the following **contributions**:

- 1) We present a new method for constructing dynamic attributed edge-colored customer networks with a labeling intuition that allows for re-influencing prediction.
- 2) We develop DGNN models by combining GNNs and RNNs, i.e., GATs and GINs with Long-Short Term Memory (LSTM) and Gated Recurrent Unit (GRU) models. We rigorously evaluate the performance of different model combinations to suggest an optimal architecture for deploying our solution.
- 3) We present a new method of inductive learning from imbalanced data by incorporating weighted loss functions, GraphSMOTE, and a rolling-window strategy in a creative way.
- 4) We thoroughly evaluate our method against baseline static GNNs and dynamic non-GNN approaches based on their predictive performance as well as ecological footprint.¹
- 5) We introduce a new profit framework for influencer prediction to decide on the optimal classification threshold and evaluate models in a profit-driven manner.

The **novelty** of our research is that we address influencer prediction on dynamic attributed edge-colored customer networks with cost-sensitive decision making using DGNNs, something that has not been done in prior research.

This paper is organized as follows. Section II provides an overview of the related studies. In Section III, the methodology is defined. The experimental setup is outlined in Section IV. We present the results of the experiments in Sections V and the profit framework in Section VI. Section VII discusses the implications of the results. Finally, Section VIII concludes the research by addressing its limitations and suggestions for future research.

II. RELATED WORK

A. GRAPH NEURAL NETWORKS

Deep learning's increasing popularity has brought GNNs to the forefront of representational learning for networks [16]. With their versatility, GNNs can solve node-, edge-, and graph-level tasks using supervised, semi-supervised, and unsupervised learning. GNNs can be categorized into several types based on their learning approach, task level, and complexity. Recurrent GNNs, the pioneers in this field, use the same set of parameters recurrently to learn high-level node

¹The source code and online appendices can be found at <https://github.com/Banking-Analytics-Lab/INFLECT>

representations [16]. Convolutional GNNs, a more advanced type, generalize the fixed grid structure of Convolutional Neural Networks (CNN) and use either spectral-based or spatial-based graph convolution to aggregate neighbor features and learn node representations [16]. Graph Autoencoders learn node embeddings in an unsupervised manner by reconstructing the adjacency matrix, while Spatial–Temporal GNNs enable dynamic network analysis by incorporating the time dimension [16].

Although there are numerous GNNs available, they share a similar design pipeline [10]. It includes the steps of identifying network structure and type, designing the loss function, and building the model using computational modules [10]. Network structure can be either explicit, where the network is given upfront, or implicit, where the network is to be built from the task at hand. The main network types are directed or undirected, homogeneous or heterogeneous, and static or dynamic networks. Loss functions are dependent on the task type, including its level and the type of supervision [10]. Computational modules can be classified into propagation, sampling, and pooling modules. A propagation module is used to propagate feature and topological information between nodes and can include convolution operators, recurrent operators, or skip connections. A sampling module can be utilized to sample large networks, while a pooling module is used to extract information from nodes [10].

B. DYNAMIC GRAPH NEURAL NETWORKS

A network that changes over time, with appearing/disappearing nodes and edges, is called a dynamic network [13]. Research shows that incorporating the temporal component of networks into the prediction task enhances the models' performance [14]. Dynamic GNNs are often represented using the encoder-decoder structure, where the encoder focuses on learning node embeddings and the decoder is used to solve prediction tasks such as node or edge classification. Zhu et al. examined different encoder-decoder architectures for supervised dynamic network learning, with techniques categorized as either Discrete Time Dynamic Graph (DTDG) or Continuous Time Dynamic Graph (CTDG) learning [17]. The difference between DTDG and CTDG learning is in how they handle time, with DTDG models using network snapshots that capture the topology at a point in time, and CTDG models considering networks as an event stream that updates over time. A further distinction can be made within DTDG based on how structural and temporal patterns are captured by the encoder-decoder composition [13]. Stacked DTDG models capture structural and temporal information in separate layers, while integrated DTDG models combine them in one layer. The former deals with combining existing layers in new ways (as in [18]), while the latter requires a design of new layers (e.g., [19]).

C. REFERRAL MARKETING

Unlike the business-to-customer (B2C) relationships leveraged by traditional marketing techniques, referral marketing

leverages customer-to-customer (C2C) relationships to promote products or services. Referral programs create additional motivation for customers to refer other potential customers in exchange for a reward. The reward can be a coupon, cash, bonus points, bonus products, etc. Targeted and referral marketing programs catalyze WOM effects by reaching influential customers to encourage them to make product recommendations [20]. Ryu and Feick found that the offer of a reward increases the likelihood of a referral, regardless of the size of the reward offered [20]. In particular, this effect is stronger for weak ties between a referrer and a potential customer, i.e., exchange relationships, as well as for weaker brands. Thus, referrals require more encouragement in the absence of strong ties or high brand recognition. However, even for stronger ties and brands, the reward should be present (non-zero) to reduce the level of inequity for existing customers. Therefore, firms need to balance the marginal revenue impact with the size of the reward and the change in the likelihood of referrals [20]. Armelini et al. found that referred customers are more loyal than non-referred customers [21]. In addition, they also found that referred customers are more valuable when the referring customer has a high customer lifetime value. This means that referral marketing is most successful when the targeted customers are not only influential, but also profitable. Also, referred customers have a longer customer lifetime than non-referred customers [7].

D. INFLUENCER DETECTION IN NETWORKS

Influencer detection has been extensively researched in the social network domain. Valuable insights into a user's network influence can be gleaned from their personal information, platform interactions, and the dynamics of user interactions. Early research focused on different approaches to detect influential network nodes, including network centrality measures, prestige ranking algorithms, and information diffusion methods [6]. Influential nodes can also be discovered by following the graph-driven branching process, which generalizes network centrality metrics [22]. Text data, such as online forums or tweets, can serve as a source for semantic analysis, which has been shown to be useful for filtering out irrelevant network interactions and identifying key nodes using classical authority discovery algorithms [23]. Considering different types of nodes and their interactions can also be advantageous for influencer detection, as in the case of the SNet model [24]: by examining relationships not only among users, but also between users and posts, influence can be estimated, and detected influencers can be leveraged for influencer marketing. In addition, the scope extends beyond different user categories to heterogeneous social connections that manifest as multi-layered networks. Combined with a heterogeneous degree ranking algorithm, they provide a powerful framework that allows for a better overview of the influential nodes [5].

Influencer detection can also borrow the insights from game theory and deep learning. The former is illustrated by the

EDGly model, which goes beyond basic centrality metrics by using the ensemble clustering for graphs to identify important partitioned communities and Shapley value games to uncover influential nodes within these communities [25]. The latter is represented by a growing adoption of GNNs in the influencer detection domain.

A recent systematic literature review by Seyfosadat and Ravanmehr [26] outlines the studies on identifying influencers in social networks. Among the machine learning-based and graph-based approaches reviewed by [26], the studies that use GNNs to detect influencers are mentioned, including the studies by [11], [27], and [28]. In addition, Graph Influence Networks have emerged as a specialized GNN technique tailored for influencer discovery by exploiting the inherent local graph structures to identify influential neighbors of a node [12]. Table 1 provides an overview of the studies mentioned above.

E. RESEARCH GAP

The current literature reviewed in section II-D has extensively studied the problem of influencer detection using various economic, statistical and machine learning techniques. However, the majority of these studies consider influencer detection in the online social media platforms and do not cover other domains such as corporate customer networks, which are dynamically changing. Moreover, the research on influencer detection with GNNs is limited because it lacks the incorporation of more powerful and recent techniques.

In our previous work [29], we introduced an influencer detection task in dynamic customer networks, which are conceptually different from social media networks, and defined influencers as customers who have at least once referred other customers, with the influencing behavior being propagated to the future: “once an influencer - always an influencer”. In this study, we reformulate the transductive influencer detection into the inductive influencer *prediction* problem for imbalanced customer networks to allow for re-influencing within the existing network nodes and generalizing towards nodes not seen by the model. We also extend the analysis by incorporating a more sophisticated GNN encoder, i.e., Graph Isomorphism Network, employing focal loss, and performing a profit-driven model evaluation, which has not been previously explored for influencer prediction.

III. METHODOLOGY

A. PROBLEM DEFINITION

Following the pipeline developed by Zhu et al. [17], we define influencer prediction as a supervised node-level learning on a dynamic heterogeneous undirected network following the stacked DTDG approach. We consider a dynamic network as an ordered sequence of T network snapshots $G = (G_1, G_2, \dots, G_T)$, $G_t = (V_t, E_t)$ where V_t is the set of nodes at time t , $|V_t| = n$, and $E_t \subseteq V_t \times V_t$ is the set of edges at time t , representing the connections between nodes, $t \in \{1, \dots, T\}$. Each node is labelled with a vector \mathbf{I}_v^t at time t with values

in $\{0, 1\}$, an influencer is labelled as 1, and vice versa. Each node $v \in V_t$ is attributed with node features \mathbf{x}_v^t at time t . Each edge $(u, v) \in E_t$ between nodes $u, v \in V_t$ is characterized by edge features $\mathbf{e}_{u,v}^t$ at time t with values in $\{0, 1\}$. For the static problem formulation (i.e., GNN specification), t is omitted. The task is node classification, i.e., predicting a label \mathbf{I}_v^{T+k} for a node $v \in V_{T+k}$ at time $T+k$.

B. GNN-RNN CONFIGURATIONS

Introduced by Veličković et al., GATs consist of layers in which nodes use attention mechanism to emphasize connections with neighbors [30]. Stacking these layers gives implicit weights to neighboring nodes without requiring full knowledge of the graph. Brody et al. proposed a more expressive version, GATv2, which employs dynamic graph attention [31]. In GATv2, the attention ranking depends on the query node, allowing all nodes to attend to each other. In a single layer of GATv2, node features as well as edge features (if any), are taken as input, and a set of node embeddings is produced as an output. To accomplish this, each node attends to its one-hop neighbors, including itself, by calculating attention coefficients $\alpha_{i,j}$ and the final node embeddings \mathbf{x}_i' as follows:

$$\alpha_{i,j} = \frac{\exp(\mathbf{a}^\top \text{LeakyReLU}(\Theta[\mathbf{x}_i \parallel \mathbf{x}_j \parallel \mathbf{e}_{i,j}]))}{\sum_{k \in \mathcal{N}(i) \cup \{i\}} \exp(\mathbf{a}^\top \text{LeakyReLU}(\Theta[\mathbf{x}_i \parallel \mathbf{x}_k \parallel \mathbf{e}_{i,k}]))} \quad (1)$$

$$\mathbf{x}_i' = \left\| \begin{array}{l} \sigma(\alpha_{i,i}^m \Theta^m \mathbf{x}_i + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j}^m \Theta^m \mathbf{x}_j) \end{array} \right. \quad (2)$$

where $\mathcal{N}(i)$ is a set of nodes adjacent to i , Θ is a weight matrix, $\alpha_{i,i}^m$ is an attention coefficient computed by the m -th attention mechanism, \mathbf{a}^\top is a weight vector of the attention mechanism, \parallel is concatenation, \mathbf{x}_i are the features of a node i and $\mathbf{e}_{i,j}$ are the features of an edge between nodes i and j .

Recent research on GNNs has focused on maximizing the representational power of learned embeddings. GINs, introduced by Xu et al. [32], have been shown to be as powerful as the Weisfeiler-Lehman graph isomorphism test (WL test) [33]. GINs are thus considered to be among the most powerful GNNs in terms of discriminative and representational power [32]. Graph isomorphism is closely related to graph learning: isomorphic graphs should be mapped to the same representation, and vice versa. To achieve the same power as the WL test, a GNN's neighbor aggregation must be injective. GINs employ a Multi-Layer Perceptron (MLP) to model and learn these injective functions [32]. We use the modified GIN operator [34] which can incorporate edge features into the aggregation procedure, contrary to the original GIN operator [32] that can only deal with node features, as follows:

$$\mathbf{x}_i' = h_\Theta((1 + \epsilon) \cdot \mathbf{x}_i + \sum_{j \in \mathcal{N}(i)} \text{ReLU}(\mathbf{x}_j + \mathbf{e}_{j,i})) \quad (3)$$

where h_Θ is an MLP, \mathbf{x}_i is a feature vector of node i , \mathbf{x}_j is a feature vector of node j from a set of nodes $\mathcal{N}(i)$ adjacent to

TABLE 1. Related work on GNNs for influencer detection.

Ref.	Description	GNN used	Dynamic?
[11]	Network topology and text message semantics are combined for influencer discovery, leveraging both a language attention network (LAN) and a GCN: LAN is used as a subject filter, while GCN is used to model influence propagation.	GCN	No
[27]	Inspired by GCN workings, the critical node identification problem in complex networks is transformed into a regression problem: considering both adjacency matrix and convolutional neural networks (CNN), a feature matrix is generated for each node, and CNN is used to train and predict the influence of nodes.	GCN-inspired	No
[28]	Based on the fixed-size neighbor networks identified by the Breadth First Search algorithm and node features derived from the network (degree, betweenness, closeness and clustering coefficient), the GCN model is used to learn latent representations of the nodes. These representations are then used to solve the downstream prediction task.	GCN	No
[12]	The most influential nodes in a graph are identified using both topological structures and node features by (1) constructing an influence set of a node based on the local graph structure; (2) measuring the feature influences of the nodes in the influence set; (3) selecting influential neighbors based on structural accessibility and task relevance.	Novel GINN	No

i , $\mathbf{e}_{j,i}$ is an edge features vector. ϵ determines the importance of the target node compared to its neighbors and is set to 0 following the findings of Xu et al. where learning ϵ yielded no gain in fitting training data [32].

While the aforementioned GNNs proved effective, they are static and cannot capture changes that occur over time. To address this, we incorporate RNNs into the decoder component of our framework, specifically the Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) models [35], [36]. The LSTM model aims to learn long-term dependencies with an LSTM cell consisting of several gates, including the input gate i_t , the forget gate f_t , the cell gate g_t , and the output gate o_t [35]. The forget gate is used to decide which information to keep and which to discard, while the input and cell gates use the input and previous hidden state to update the cell state. The cell state c_t can then be updated with the outputs of the input, cell, and forget gates as follows:

$$\begin{aligned}
i_t &= \sigma(W_{ii}\mathbf{x}^t + b_{ii} + W_{hi}h_{t-1} + b_{hi}) \\
f_t &= \sigma(W_{if}\mathbf{x}^t + b_{if} + W_{hf}h_{t-1} + b_{hf}) \\
o_t &= \sigma(W_{io}\mathbf{x}^t + b_{io} + W_{ho}h_{t-1} + b_{ho}) \\
c_t &= f_t \odot c_{t-1} + i_t \odot \tanh(W_{ig}\mathbf{x}^t + b_{ig} + W_{hg}h_{t-1} + b_{hg}) \\
h_t &= o_t \odot \tanh(c_t)
\end{aligned} \quad (4)$$

where \mathbf{x}^t is the output embeddings of a GNN model at time t , h_{t-1} is the hidden state at time $t - 1$, n_t is the new gate used to update a hidden state using a reset gate, W are learnable weights, b are bias variables, σ is the sigmoid function, \tanh is the hyperbolic tangent function, \odot is the Hadamard product.

Finally, the output gate, together with the cell state, is used to update the hidden state received from the previous timestamp $t - 1$, as can be seen from (4). These gates allow the cell state, which can be thought of as the long-term memory of the network, to be modified at each timestamp t . By leveraging these operations, the LSTM can effectively capture and learn from long-term patterns.

The GRU model, which is a simpler version of the LSTM, consists of two primary gates: the reset gate r_t and the update gate z_t . The update gate operates similarly to the input gate in LSTM models, and the reset gate functions in a manner comparable to the forget gate. Unlike the LSTM, the GRU does not have a cell state and stores long-term memory directly

in the hidden states. Each hidden state is updated with the previous hidden state values and the new gate value obtained from the reset gate output as follows:

$$\begin{aligned}
r_t &= \sigma(W_{ir}\mathbf{x}^t + b_{ir} + W_{hr}h_{t-1} + b_{hr}) \\
z_t &= \sigma(W_{iz}\mathbf{x}^t + b_{iz} + W_{hz}h_{t-1} + b_{hz}) \\
n_t &= \tanh(W_{in}\mathbf{x}^t + b_{in} + r_t \odot (W_{hn}h_{t-1} + b_{hn})) \\
h_t &= (1 - z_t) \odot n_t + z_t \odot h_{t-1}
\end{aligned} \quad (5)$$

where \mathbf{x}^t is the output embeddings of a GNN model at time t , h_{t-1} is the hidden state at time $t - 1$, n_t is the new gate used to update a hidden state using a reset gate, W are learnable weights, b are bias variables, σ is the sigmoid function, \tanh is the hyperbolic tangent function, \odot is the Hadamard product.

As can be seen from (4) and (5), the input to both the LSTM and GRU models are the output embeddings \mathbf{x}_i^t of a GNN encoder produced at time t for each node $i \in V_t$, with the subscript i omitted for simplicity.

To obtain final predictions, we require an additional Fully Connected Network (FCN) on top of the GNN-RNN configuration. The FCN consists of two linear layers followed by ReLU and sigmoid activation functions, respectively. It takes as input the node embeddings, i.e., the hidden states of the last output layer of the RNN model, and outputs the probability of a node being an influencer. The architecture is shown in Figure 1.

C. BASELINE CONFIGURATIONS

a: FEATURES(+PR)-RNN

Following the strategy of Öskarsdóttir et al. [37], we can enrich non-relational classifiers with network features. One such network feature that can summarize the importance of a node is PageRank (PR) [38]. Originally developed to measure the importance of web pages, PR can be used to measure the importance of nodes based on the number of incoming edges and their importance. The PageRank measure and its variants have been shown to boost predictive models [2]. The non-personalized PR values can be calculated for each node as $PR(u) = (1 - d) + d \sum_{v \in B_u} \frac{PR(v)}{C(v)}$ where B_u is a set of nodes linked to u , d is a damping factor set to 0.85, $C(v)$ is an out-degree of a node v .

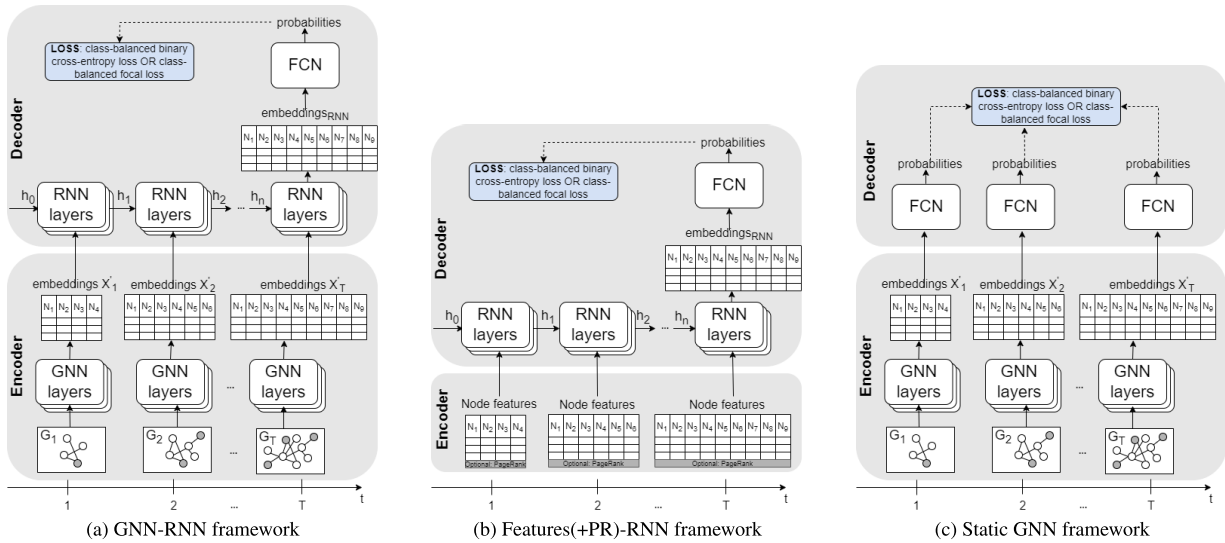


FIGURE 1. Model architectures. (a) For each time point, the model takes the network as input, i.e., network connectivity, node & edge features. This input is first passed through GNN layer(s), which generate network embeddings in Euclidean space. These embeddings, together with the hidden state from the previous step (initially a tensor of ones), are then fed into RNN layer(s). This process is repeated for each time point. At the last time point, the output of the RNN model is fed into a fully connected layer to produce the final probabilities of a node being an influencer. The weights of these layers are learned by optimizing a loss function, using either class-balanced binary cross-entropy loss or class-balanced focal loss (see Section IV-C). (b) For each time point, the model takes node features as input (with the PageRank feature added as well in the corresponding configurations). These features, along with the hidden state from the previous step (initially a tensor of ones), are then fed into the RNN layer(s). This process is repeated for each time point. At the final time point, the output of the RNN model is fed into a fully connected layer to produce the final probabilities of a node being an influencer. The weights of these layers are learned by optimizing a loss function, using either class-balanced binary cross-entropy loss or class-balanced focal loss (see Section IV-C). (c) For each time point, the model takes the network as input, i.e., network connectivity, node & edge features. This input is passed through GNN layer(s) that generate network embeddings in Euclidean space. The embeddings are then passed to a Fully Connected Layer to produce the final probabilities of a node being an influencer. The weights of these layers are learned by optimizing a loss function, using either class-balanced binary cross-entropy loss or class-balanced focal loss (see Section IV-C). There is no temporal component in this configuration, and time points are used in a batch manner.

Since the PR value represents the relative importance of a node within one component, we calculate its value separately within each of the components and scale the value with the size of the component. Since the network is attributed, PR values can be seen as an additional feature of the node, and this enriched feature set can be seen as an encoder part and can be subsequently used for dynamic node classification with RNNs. Static node features together with PR are used as input to the RNN models, i.e., LSTM and GRU models. The aforementioned architecture is shown in Figure 1b.

b: STATIC GNN

This configuration consists of a GNN encoder similar to the GNN-RNN configuration described in Section III-B, while its decoder contains only an FCN network that is fed with the embeddings produced by the GNN model. For the GNN encoder, we use GAT and GIN models described in Section III-B. The aforementioned architecture is shown in Figure 1c.

IV. EXPERIMENTAL SETUP

A. NETWORKS

The data used in this study is sourced from a Super-App company operating in Latin America that offers both a delivery app and credit card services. The data covers a period of nine months, during which aggregated information on delivery app usage was used to create monthly snapshots of the network,

while credit card usage data was used to construct node features (see Appendix A, Table 6). Numerical features are normalized using min-max normalization to stabilize learning and speed up convergence. Categorical features are one-hot encoded. To account for new nodes appearing in the network over time, we add artificial node features for the future nodes (not used in the backpropagation process during training) with features that are zeroed out. The network only includes credit card customers who use the delivery app, with three types of connections based on the app usage, as described in Table 2. To account for these types of connections, edge features are one-hot encoded with information about the connection type(s), allowing multiple connection types to be represented per edge. The network is dynamic, with new nodes and edges appearing over time, while existing nodes and connections persist, leading to its continuous growth.

TABLE 2. Edge types.

Type	Description
Credit card	There exists an edge between customers who used the same credit card in the delivery app.
Geohash	There exists an edge between customers who ordered more than 4 times in the close geographical proximity (Geohash 7–152.9m x 152.4m).
Contact	There exists an edge between customers if one of them has the other one in their phone contacts.

Our analysis is based on data from three different cities, each represented by a separate network (network characteristics

are provided in Appendix B, Figure 7). The cities are diverse in terms of population size and income, representing both middle-income and high-income cities, with Purchase Power Parity-adjusted Gross Domestic Products (PPP GDP) per capita ranging from \$10,000 to \$35,000.² The networks vary in size, with City 1 having the smallest network and City 3 having the largest. Over time, all three networks have grown in terms of both the number of nodes and the number of edges. However, the data is heavily imbalanced as the proportion of influencers among the nodes is relatively low, ranging from about 1.5% in the last month to about 5% in the first month. To address this issue, we use the synthetic node generation method of GraphSMOTE, which performs oversampling on top of the embedding space constructed by the feature extractor [15]. The original GraphSMOTE constructs an embedding space to encode the similarity between the nodes, and uses this space to generate new samples, while also generating edges to model the relationship information for these new samples [15]. We use a part of GraphSMOTE that interpolates the generated embeddings of the minority class to create new synthetic embeddings of that class. In particular, we oversample the final embeddings generated by the RNN model for the GNN-RNN model configurations or by the GNN encoder in static GNN configurations with an oversampling scale of 0.5, as scales below 0.8 typically lead to better GNN performance [15]. In particular, we have included a positive oversampling ratio as a hyperparameter specification, while also allowing for a non-oversampling strategy with a ratio of 0.

B. INFLUENCER DEFINITION

An influencer is defined as a super-app user with a credit card who successfully refers another super-app user without a credit card, resulting in the referred person obtaining a credit card. In particular, we use the following intuition for node labeling. Node labels are determined based on referral information from credit card applications. A referrer is always a node in the network, while the referred person may or may not be in the network, depending on when they obtain a credit card. A node is labeled as an influencer at the time of the actual referral and remains labeled as such until the referred person obtains a credit card, which is the point at which the company has access to their usage data, up to a maximum of six months. At that point, the influencer label is changed to a non-influencer label, unless the customer has made other referrals in the past six months, as obtaining a credit card is no longer influenced solely by a referral, but by other factors as well, and the information asymmetry between referred and non-referred customers vanishes over time [39]. To maintain the connectivity expressed by a referral, an edge of type ‘‘Contact’’ is created between the referrer and the referred person by adding the referred person as a node with an edge to the referrer. Figure 2 illustrates the labeling process. We do not distinguish between the influencing strength (e.g., whether

one node is considered more influential than another one) as we are dealing with binary classification.

C. LOSS FUNCTIONS

As part of our model optimization, we explore two loss function configurations: weighted binary cross-entropy and class-balanced focal loss. A traditional binary cross-entropy (BCE) loss is not suitable for handling class imbalance, since it treats both classes equally and prioritizes overall accuracy rather than class-specific performance. Therefore, a custom binary cross-entropy loss function is needed to handle class imbalance and weight minority examples accordingly to balance the trade-off between recall and precision, as follows:

$$\ell_c(x, y) = \sum_{i=1}^n (-w_{n,c} [p_c y_{n,c} \cdot \log(\sigma(x_{n,c})) + (1 - y_{n,c}) \cdot \log(1 - \sigma(x_{n,c}))]) \quad (6)$$

where $c = 1$ for single-label binary classification, n is the number of samples, $p_c = \frac{n_{neg}}{n_{pos}}$ is the weight of the positive class, n_{neg} (n_{pos}) is the number of samples in the negative (positive) class, σ is a sigmoid function.

To go beyond simple inverse sample weighting, Cui et al. [40] proposed a class-balanced focal loss, which addresses the challenge of class imbalance in a more sophisticated way by considering both the effective number of samples and the relative loss for well-classified samples [40], [41], as follows:

$$\ell_c(x, y) = \sum_{i=1}^n \left(-\frac{1 - \beta}{1 - \beta^{n_c}} \sum_{i=1}^C (1 - \sigma(x_{n,c}))^\gamma \log(\sigma(x_{n,c})) \right) \quad (7)$$

where C is the total number of classes, n_c is the number of samples in the ground-truth class c , $\beta \in [0, 1)$ is a hyperparameter from the effective number of samples for class term, γ is a focusing hyperparameter.

By adjusting a class-balanced term between no reweighting and reweighting by inverse class frequency, the class-balanced focal loss is able to handle class imbalance more effectively [40].

D. HYPERPARAMETER TUNING

The best model configuration is found by an exhaustive grid search over the hyperparameter space (Table 3). We select the best model based on the average AUPRC value calculated on seen and unseen nodes from the validation set as well as based on stability of the training process, i.e., we ensure that the model converges. In addition, the number of heads in GAT is fixed at 2 to keep the computational complexity manageable. We also set the focal loss hyperparameters (see Section IV-C) γ and β to 2 and 0.999, respectively, as these values have been found to perform best [40], [41]. For both the GNN-RNN and Static GNN models, we applied two regularization techniques that are commonly used for GNNs to improve their performance and generalization capabilities: layer normalization and a 50% dropout rate.

²For privacy reasons, we cannot disclose the names of the cities.

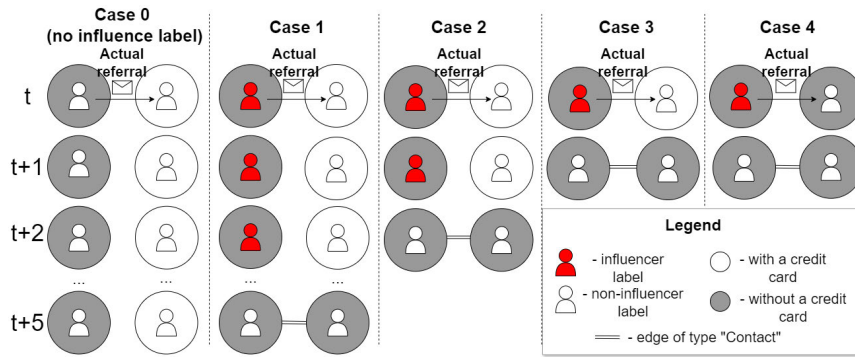
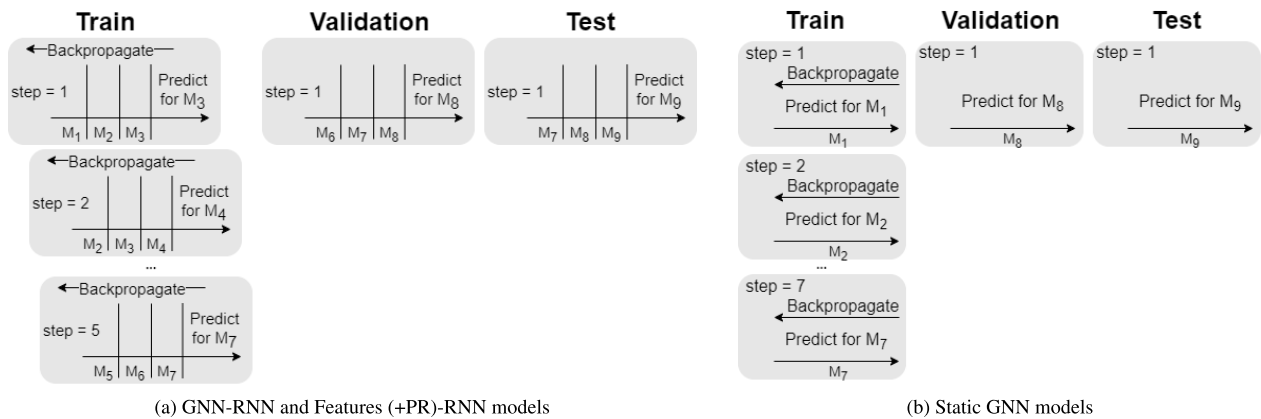


FIGURE 2. Labelling intuition.



(a) GNN-RNN and Features (+PR)-RNN models

(b) Static GNN models

FIGURE 3. Train/validation/test setup.

TABLE 3. Hyperparameters tuning.³ The models are trained on NVidia A100 GPU for 300 epochs with a learning rate of 0.0001 using the ADAM optimizer [42]. The models are implemented using Pytorch Geometric [43].

Hyperparameter	Range	Model
GNN hidden dimensions, GNN embeddings	256, 512	GNN-RNN, Static GNN
GNN layers	2, 4	GNN-RNN, Static GNN
RNN hidden dimensions	256, 512	GNN-RNN, Features(+PR)-RNN
RNN layers	1, 2	Features(+PR)-RNN
SMOTE sample rate	0, 0.5	All models
Loss function	BCE Loss, Focal Loss	All models

E. GNN-RNN MODELS TRAINING

To evaluate model performance and ensure that it can generalize, we use a rolling window strategy with a window size of three months and a shift of one month to split the data into train, validation, and test subsets according to the out-of-time validation and testing strategy (Figure 3).

The train data spans seven months, resulting in five three-month windows. Within each window, the trained model is

³For City 3, GAT-RNN models, the following configurations are not feasible due to computing limitations: all models with GNN layers = 4; configurations with GNN hidden dimensions = RNN hidden dimensions = GNN embeddings = 512.

used to make predictions for the last month. We use a network snapshot of each month as input to the GNN encoder, which produces embeddings that are subsequently fed to the RNN decoder, with the input hidden states initialized as a tensor of ones for the first time window. For all subsequent time windows, the hidden state produced by the first month of the previous time window is used as the input hidden state. The output produced by the RNN decoder at time t is used as the input hidden state at time $t + 1$. At the final timestamp of the window, the hidden states produced by the RNN decoder are fed into the FCN decoder to generate final probabilities, and backpropagation occurs during training.

Both validation and test data consist of one three-month window. The trained models are applied to a validation set that is used to measure model generalization capability by checking its convergence during training by inspecting the validation loss. Additionally, a validation set is used to select the best hyperparameter configuration (see Section IV-D). A test set is used for the final model comparison (with the best hyperparameter setting chosen on a validation set).

F. BASELINE MODELS TRAINING

The process of splitting the data into train, validation, and test subsets and the training methodology for the Features(+PR)-RNN models are similar to those of the GNN-RNN models, with the only difference being the encoder part.

In this case, instead of using GNN embeddings, the encoder consists only of node features (including the PageRank feature for PR configurations), which are fed directly into the RNN decoder. The rest of the training process follows the steps outlined in Section IV-E (Figure 3).

Unlike the previous models with the RNN component, static GNN models do not take dynamics into account. The training data for the static model spans seven months, with predictions made for each month and the training performance averaged over those months. The validation and test sets consist of data from a single month each, with predictions made for their respective months (Figure 3b).

G. PERFORMANCE METRICS

We are solving the influencer prediction problem in a profit-driven manner, and therefore the final model threshold decision is made based on profit rather than predictive performance at that particular threshold. Therefore, we use threshold-independent metrics.

The area under the receiver operating characteristic (ROC) curve (AUC), which plots the true positive rate (TPR) against the false positive rate (FPR) for different classification thresholds, is often calculated to compare model performance, with a range of values from 0 to 1 and a baseline of 0.5. In the presence of class imbalance, AUC may not be a reliable performance indicator as it is necessary to consider the skewed class distribution. The Precision-Recall (PR) curve is a better alternative for imbalanced data as it summarizes performance over different thresholds with precision and recall on its axes [44]. To facilitate model comparison, the area under the PR curve (AUPRC) can be calculated with the baseline corresponding to the fraction of positives. While the AUPRC metric is our primary criterion for evaluating performance, we also report the AUC metric.

For both AUC and AUPRC, we report the values separately for seen and unseen nodes. The term “seen nodes” refers to nodes that the model encountered during backpropagation in training, whereas “unseen nodes” refers to nodes that only appear in the network later. Hence, “unseen nodes” for the validation data are the nodes that appeared in the last month of the validation window, while the “unseen nodes” for the test data include new nodes that appeared in the last month of the test window as well as in the last month of the validation window, as all of them were not part of the training. Model performance on seen nodes shows the models’ applicability for transductive learning which focuses on making predictions specifically for the given training instances without generalizing to unseen instances, while performance on unseen nodes highlights the ability to generalize from observed data to make predictions on unseen instances, i.e., inductive learning.

We also report training time and carbon footprint estimate measured in equivalent grams of CO₂ (*gCO₂eq*), as we also want to evaluate the tradeoff between energy efficiency and prediction performance [45].

V. RESULTS

Table 4 displays the values of the performance metrics. The results for City 1 demonstrate that GIN-RNN models perform similarly to dynamic non-GNN models in terms of AUC for both seen and unseen nodes, with GIN-GRU being the best performing model for seen nodes. The AUPRC values for seen nodes are also comparable across these model types, with GAT-LSTM achieving the highest value. Notably, the Static GIN model delivers exceptional performance in terms of AUC on unseen nodes. Nevertheless, the highest AUPRC values for unseen nodes are obtained by the GIN-LSTM and GIN-GRU models. The GIN-RNN models achieve comparable time performance to the baseline models, while the GAT-RNN models have the longest training time and produce the highest amount of CO₂ emissions due to an expensive attention computation.

TABLE 4. Models performance on test data. Best model configurations can be found in Tables 7-9.

Model	AUC seen nodes	AUC unseen nodes	AUPRC seen nodes	AUPRC unseen nodes	Total time, s.	CO ₂ , g.
City 1						
GIN-LSTM	0.75±0.02	0.67±0.06	0.08±0.02	0.11±0.06	57	5.21
GIN-GRU	0.79±0.02	0.63±0.06	0.09±0.02	0.10±0.06	48	3.68
GAT-LSTM	0.70±0.03	0.60±0.07	0.12±0.03	0.07±0.04	188	16.29
GAT-GRU	0.65±0.03	0.55±0.06	0.08±0.02	0.05±0.04	183	17.12
Features(+PR)-LSTM	0.77±0.02	0.68±0.06	0.10±0.03	0.08±0.05	45	3.63
Features(+PR)-GRU	0.77±0.02	0.67±0.06	0.09±0.02	0.09±0.05	41	3.01
Features-LSTM	0.75±0.02	0.65±0.07	0.09±0.02	0.08±0.05	90	5.88
Features-GRU	0.77±0.02	0.68±0.06	0.09±0.02	0.09±0.05	43	2.75
Static GIN	0.69±0.03	0.72±0.06	0.05±0.01	0.09±0.05	36	2.21
Static GAT	0.64±0.03	0.54±0.06	0.03±0.004	0.03±0.02	94	9.27
City 2						
GIN-LSTM	0.71±0.02	0.78±0.04	0.08±0.02	0.19±0.08	54	3.90
GIN-GRU	0.67±0.02	0.80±0.04	0.07±0.02	0.22±0.08	61	4.31
GAT-LSTM	0.68±0.03	0.74±0.05	0.06±0.01	0.16±0.07	104	9.21
GAT-GRU	0.66±0.03	0.73±0.06	0.06±0.02	0.18±0.08	202	18.82
Features(+PR)-LSTM	0.70±0.02	0.64±0.07	0.04±0.01	0.07±0.04	47	3.54
Features(+PR)-GRU	0.71±0.02	0.65±0.07	0.06±0.02	0.07±0.03	38	2.33
Features-LSTM	0.70±0.02	0.64±0.07	0.05±0.01	0.07±0.03	48	3.82
Features-GRU	0.71±0.02	0.66±0.07	0.05±0.01	0.08±0.04	99	5.86
Static GIN	0.66±0.03	0.83±0.03	0.05±0.01	0.13±0.06	37	2.20
Static GAT	0.60±0.03	0.70±0.05	0.03±0.003	0.07±0.04	60	4.83
City 3						
GIN-LSTM	0.73±0.01	0.78±0.02	0.08±0.01	0.13±0.03	517	40.61
GIN-GRU	0.75±0.01	0.78±0.02	0.09±0.01	0.13±0.02	622	57.51
GAT-LSTM	0.69±0.01	0.68±0.02	0.09±0.01	0.12±0.02	685	70.90
GAT-GRU	0.68±0.01	0.70±0.02	0.09±0.01	0.13±0.03	525	45.06
Features(+PR)-LSTM	0.70±0.01	0.67±0.02	0.03±0.002	0.03±0.006	433	34.71
Features(+PR)-GRU	0.71±0.01	0.66±0.02	0.04±0.003	0.05±0.01	5106	302.56
Features-LSTM	0.69±0.01	0.66±0.02	0.03±0.002	0.03±0.005	2823	166.58
Features-GRU	0.71±0.01	0.67±0.02	0.04±0.002	0.04±0.01	397	25.44
Static GIN	0.69±0.01	0.72±0.02	0.05±0.004	0.10±0.02	360	22.87
Static GAT	0.63±0.01	0.61±0.02	0.04±0.003	0.12±0.03	909	81.86

For the larger network of City 2, the GNN-RNN models obtain similar AUC results as the dynamic non-GNN models on seen nodes. However, the best AUC values on unseen nodes are obtained by the static GIN model, followed by the GIN-GRU model. The GIN-RNN models show the best AUPRC performance on both seen and unseen nodes, with a strong dominance on unseen nodes and an AUPRC uplift of

about 0.12 from the best performing dynamic non-GNN model. Overall, the GIN-GRU model provides the best performance based on most of the metrics. In terms of time and energy performance, the pattern is the same as for City 1: complex GAT-RNN models take the longest time to train and produce maximum carbon emissions.

For the largest network of City 3, it becomes apparent that the GNN-RNN models outperform the baseline models. In particular, the GIN-RNN models achieve the highest AUC values on both seen and unseen nodes, indicating their superior performance. Additionally, in terms of AUPRC on seen and unseen nodes, all GNN-RNN models perform exceptionally well, with no clear winner among them.

In summary, more sophisticated models, and in particular the GIN-GRU combination, perform better as the network grows, becoming dominant in the largest network of all. This is consistent with other results in deep learning, that confirm how data-hungry these models can be: in order to improve on simpler alternatives, it is necessary to have enough data diversity to achieve effective results. In terms of computational time, models with pure attention mechanisms are the most complex to train, which is consistent with their quadratic complexity. In our models, this does not translate into better performance, so modelers should carefully test other alternatives to benchmark the right model. Interestingly, the use of a balanced error measure, AUPRC, shows that more powerful models do present better results. As discussed earlier, the AUC performance can be misleading in the presence of high imbalance, so measures such as AUPRC can be more helpful.

Figures 4a, 4b, and 4c display train/validation losses, AUC, and AUPRC convergence plots for the consistently well-performing GIN-GRU models for all three cities. These plots demonstrate that the GIN-GRU model achieves convergence across all networks. While the validation AUC grows consistently with epochs for all cities, Figure 4c reveals that the model training for the largest city, City 3, is the most stable, with AUPRC gradually increasing with each epoch. In contrast, the validation AUPRC for City 1 is the least stable. Therefore, providing more data to the model improves both its training stability and accuracy. The convergence plots for the other models can be found in Online Appendix A.

Appendix C, Tables 7-9 present the optimal hyperparameter configurations for the GNN-RNN models and the benchmark models. Across all cities, focal loss is generally the preferred loss function for most models. The optimal number of layers in the GNN encoder shows no clear pattern for City 1, while for Cities 2 and 3, shallower networks with two layers are favored. A similar trend is observed for the RNN model: shallower networks with one layer are preferred for Cities 1 and 3, whereas no distinct pattern is evident for City 2.

There is no consistent preference for the number of hidden dimensions and embeddings in the GNN encoder, although configurations with 512 hidden dimensions are slightly favored for GNN-RNN architectures. The optimal number of hidden dimensions for the RNN component increases

with data size: 256 dimensions are preferred for City 1, and 512 dimensions for City 3. Regarding the SMOTE rate, no oversampling typically results in better performance for most models. Notably, none of the GNN-RNN models benefit from oversampling, with zero rate chosen for all GNN-RNN combinations.

The analysis indicates that optimal model configurations are highly dependent on the underlying data and its specific characteristics. Although some hyperparameters, such as focal loss and the absence of oversampling, remain consistently optimal across the cities, most hyperparameters vary between the models built for different cities. Therefore, each city requires a tailored model that is trained and optimized based on its own data, as we have not been able to establish the generalizability of model configurations across cities.

VI. PROFIT EVALUATION

To make use of the model predictions, the decision on the classification threshold must be made in order to return the final predicted labels used for a marketing campaign. When dealing with severe class imbalance, the default threshold of 0.5 used for the balanced tasks can lead to poor performance [46]. Therefore, we tune this threshold by applying the predict-then-optimize approach introduced by Vanderschueren et al. [8]: a cost-insensitive influencer detection model is used to optimize decision making, i.e., tuning a classification threshold to obtain the maximum expected profit of potential customers brought by influencers. Since profit is customer-specific, we use instance-dependent cost-sensitive thresholding [8].

We assume that the influencer detection model is used for targeted marketing, i.e., the influencers predicted by the model are contacted with a contact cost c . True influencers are also rewarded at a cost of ic . As we know from the data provider, all the customers have been targeted during a time frame of the study. However, the assumption that this is permanent is unrealistic, so we need to model the impact of a targeted policy. Hence, we can introduce a probability p of not referring new customers while having the opportunity to do so (being an influencer). Since we know that everyone has been targeted, true influencers have a probability of not referring equal to 0 ($p = 0$). For non-influencers, this probability is a free parameter to be estimated ($p \neq 0$). Consequently, $(1 - p)$ represents the probability that an actual influencer will still bring new customers even without being targeted.

The total profit of the campaign depends on the accuracy of the model's predictions and comprises the profit of four different categories, each constructed based on the actual and predicted influencer labels (Table 5). The first category is the Target Influencer (TI) category, which corresponds to the True Positive predictions. The profit of this category is calculated as the future profit of potential customers P_u referred by an influencer u , reduced by the incentive cost ic . In addition, a contact cost c is subtracted once, as the action to contact

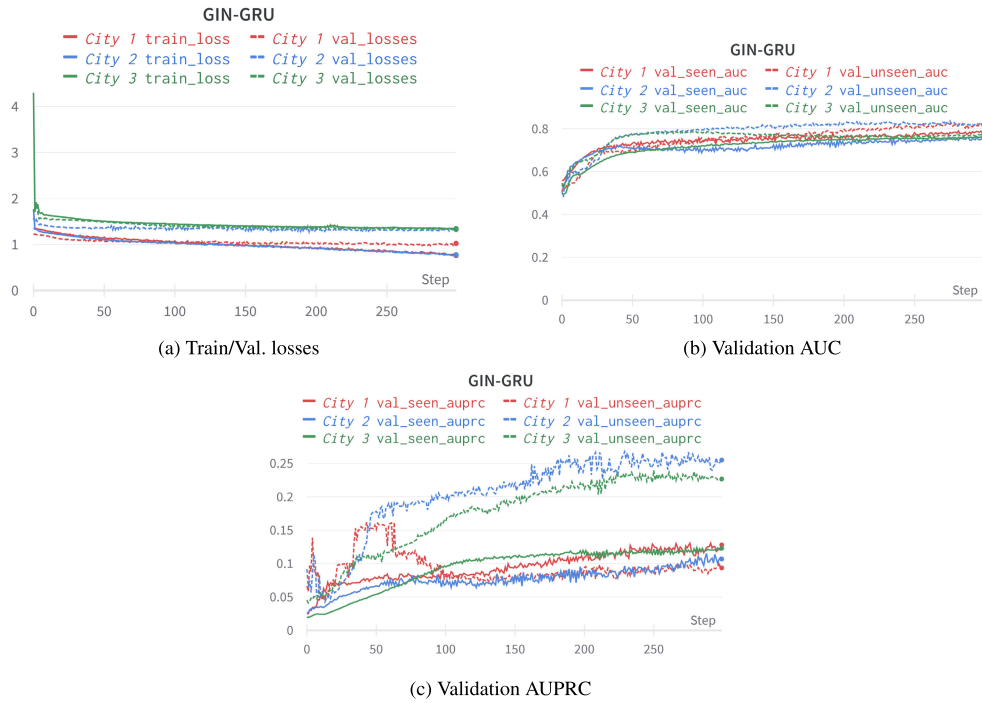


FIGURE 4. Convergence plots, GIN-GRU models.

TABLE 5. Profit. N_u is the set of people referred by u , c is a contact cost, ic is a reward cost, P_v is a future 6-month profit of a customer v calculated as an interest earned on the outstanding balance on a credit card; p is a probability of not referring new customers while having the opportunity to do so (being an influencer).

		Actual label	
		Influencer	Non-influencer
Predicted label	Influencer	$P_u(TI) = \sum_{v \in N_u} (P_v - ic) - c$	$P_u(TNI) = -c$
	Non-influencer	$P_u(NTI) = (1 - p) \sum_{v \in N_u} (P_v - ic)$	$P_u(NTNI) = 0$

the customer may have a cost.⁴ The *TNI* category represents False Positive predictions: the customers who were targeted but did not refer other customers. The profit from this category is negative because we spend money to contact them and do not get new customers in return. The *NTI* category consists of False Negatives: true influencers who were not contacted. We expect to get profit from this category of customers with the aforementioned probability $1 - p$, minus the reward cost of ic . The last category *NTNI* represents True Negative predictions that bring neither profit nor cost. We search for an optimal threshold that maximizes profit on the validation data, and then use this threshold on the test data to obtain a profit estimate. Hence, we solve the optimization problem as follows:

$$\arg \max_t P_{total}(t) = \{t \mid P_{total}(t) = \max_{t'} P(t')\} \quad (8)$$

where t is a classification threshold, P_{total} is a total profit calculated as $P_{total} = \sum_{S \in G} \sum_{u \in S} P_u$, $G = \{TI, TNI, NTNI, NTI\}$.

⁴In an app-based contact strategy, this cost can be very marginal. In traditional marketing, this cost can escalate quickly. Users must calibrate these values based on their situation.

We illustrate the aforementioned framework using the best performing GIN-GRU model for City 2. The P_{total} calculation has three parameters ic , c , and p that influence the resulting profit. We obtained a reward parameter from the data provider and normalized it to a generic value of 100 monetary units. All values are expressed based on these normalized results. We solve the optimization problem for each combination of c and p with $c \in [0, 80]$ with a step of 10 (the step is reduced to 5 in the interval $[0, 30]$ to better illustrate smaller values of c) and $p \in [0, 1]$ with a step of 0.05. We vary a threshold from 0 to the maximum probability score returned by the model, which is 0.78 for the validation month data.

Figure 5a shows how the optimal classification threshold changes for different combinations of c and p , while Figure 5b shows how these thresholds are applied to obtain a profit estimate for these combinations. In particular, when the incentive comes at no cost, the most profitable solution is to contact everyone in the customer base, regardless of the value of p . This is also the case when the maximum profit is achieved. This is not surprising, since targeting false positives (*TNI* category) comes at no cost, while targeting true positives (*TI* category) brings non-negative profit. However, this case is not realistic because of the need to contact customers through

different channels, which has a non-negative cost. The profit is also maximal when the probability of referral without being targeted is one, i.e., $1 - p = 1$, which represents a situation where no one is targeted because the profit obtained from the true influencers is higher than when the target campaign is activated, which comes with positive contact costs.

With an increasing contact cost c and a decreasing probability of referral (i.e., increasing probability p of not referring), the optimal threshold increases: it is no longer economically feasible to contact everyone. This is a more realistic situation, where the model brings added value in selecting the best customers to contact. In this case, the profit becomes smaller with increasing values of probability p . This is expected because as probability p increases, the need for the marketing campaign diminishes, as the likelihood of recommendation is high enough without the need to be stimulated by a target. As soon as the contact cost c becomes higher than 60, it is no longer feasible to target customers: the cost becomes higher than the added value of using the model's predictions, regardless of the value of p . At this point, the profit is negative and targeted marketing is no longer economically rational. Thus, profit analysis can be used to determine the optimal value of a contact cost at which the marketing campaign still makes sense from a profit perspective.

A similar analysis can be performed for the case where ic is varied instead of c (Figure 6). In this scenario, we can set c to 3^5 with $ic \in [0, 250]$ with a step of 50 and $p \in [0, 1]$ with a step of 0.05. Similar to the p vs. c case, if the probability of referring without being targeted ($1 - p$) is 1, the targeting campaign makes no sense: all the influencers will refer to other customers anyway, without being nudged to do so. For increasing probability p and decreasing reward cost ic , the optimal threshold value decreases as more customers are predicted as influencers. When the reward cost becomes too high (≥ 200), it is again optimal to not use a model, and rely on referrals without a targeting campaign. When these thresholds are applied to test data to obtain profit estimates, we observe the patterns similar to the p vs. c case: the profit decreases with increasing values of p and increasing reward costs, and becomes negative for reward costs greater than 200.

The above examples illustrate that prior knowledge is required to deploy influencer prediction models in production. The optimal thresholds should be estimated based on this prior information to get the maximum benefit from the model. The probability p , determined by personal brand attitudes and marketing preferences [20], is beyond the company's control and can be modeled in order to estimate its values depending on aforementioned characteristics. We believe this is an interesting topic for future research.

⁵The value is chosen based on the conversations with the data provider; it must be chosen based on the use case and the cost structure.

VII. DISCUSSION

As illustrated in Section V, the GNN-RNN models, in particular the GIN-GRU combination, dominate the baseline models, and this dominance is most pronounced on unseen nodes being crucial for inductive learning that focuses on generalizing to data not encountered by a model. For this configuration, the best combination of performance on seen and unseen nodes is obtained on the largest network of City 3, with consistently high values for all performance metrics. In other words, as task complexity increases, which is the case for inductive learning, the applicability of complex GNN-RNN models is justified. On the contrary, the performance of non-GNN dynamic baseline models, i.e., Features(+PR)-RNN models, deteriorates with increasing network size. While on the smallest network of City 1, these models perform comparable to the GNN-RNN models, the opposite is true for the largest network of City 3. Therefore, for smaller networks, it is advisable to choose simpler models due to the limited performance gain achieved by using GNN-RNN models, but as the networks scale, it is necessary to leverage the complexity of deep learners.

When comparing the two types of GNNs, the GIN encoder emerges as the clear winner in both GNN-RNN and static GNN model types due to its ability to be trained in less time and with fewer CO_2 emissions, while consistently outperforming the GAT in terms of detection performance. This superiority can be attributed to the complex network structures we are dealing with: GINs are universal and invariant to the graph structure, whereas GATs are more sensitive to it, which is disadvantageous for complex dynamic networks. This illustrates a trade-off between complexity and performance: the increased complexity of GATs is not justified by a relatively small performance gain. Moreover, in the realm of inductive learning, the importance of the GNN component surpasses that of network dynamics. This is particularly true in the context of the GIN encoder, where models incorporating a GNN encoder consistently exhibit superior AUC and AUPRC values compared to non-GNN dynamic models. Nonetheless, network dynamics play a crucial role and must be captured: the best results are obtained when GNNs and RNNs are coupled as a single model. Given the time and CO_2 emissions of the GIN-RNN models, they are also the best candidates for production use: on average, they are the fastest to train with the least carbon footprint. Given the rising cost of energy, this is an important advantage of these models.

The profit framework highlights the targeting advantages that the proposed models, particularly the GIN-GRU model, can provide. The experiments illustrate that the applicability of the model depends heavily on both the costs of contacting a customer and the value of the reward offered to the referrer. According to Ryu and Feick, this reward should be non-zero, and the model is found to be useful in cases with lower contact/reward costs and higher probabilities of not referring without being targeted [20]. Since the company can only

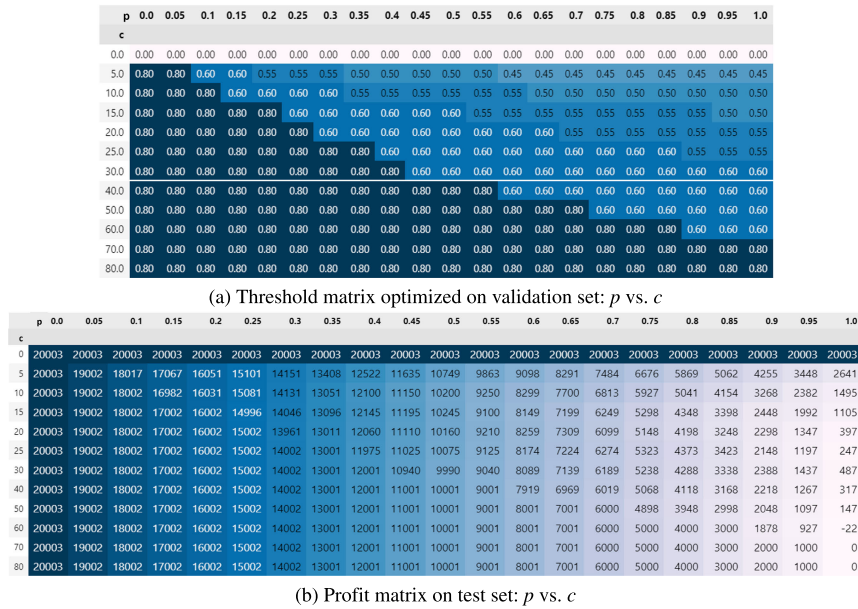


FIGURE 5. p vs. c analysis.

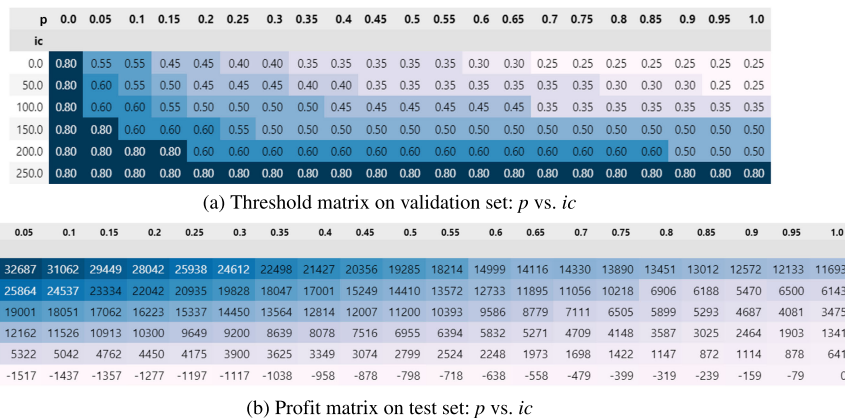


FIGURE 6. p vs. ic analysis.

influence the contact/reward costs, they should be kept as low as possible. However, since there is a positive relationship between a reward size and a referral likelihood [20], a trade-off between the two should also be considered (see future research directions in Section VIII).

VIII. CONCLUSION

Network data is getting increased attention due to its ability to reveal complex relationships that may go unnoticed by traditional data sources. Marketing is no exception, as network information can be used to promote products and services through marketing initiatives. Given that DGNNs have demonstrated effective learning from evolving network data [17], this paper investigated the optimal strategy for deploying DGNNs to solve the inductive influencer prediction problem, emphasizing profit-sensitive decision-making optimization. Four DGNN configurations were trained and evaluated on network data from three cities and compared

with dynamic non-GNN and static GNN baseline models. The evaluation was performed based on the AUC and AUPRC metrics calculated on seen and unseen nodes, as well as on the time performance and environmental impact of a model. A predict-then-optimize approach was used to illustrate the usability of the models in real business environments.

First, our study shows that using DGNNs for inductive influencer prediction is highly advantageous. By incorporating dynamic network information, these models demonstrate better generalization to unseen data while maintaining strong detection performance on nodes encountered during training. Among the different types of encoders tested, the GIN model is the most effective due to its ability to handle universal graph structures. We therefore recommend the use of DGNNs with a GIN encoder for influencer prediction tasks that require generalization to new nodes.

Next, we illustrate that applying DGNNs to larger networks provides the best combination of detection performance

and model stability. In these cases, DGNNs are particularly useful because they can effectively capture the increasing complexity that cannot be fully captured by centrality metrics such as PageRank. However, for smaller networks with less informative topology, simpler dynamic non-GNN models can perform comparably to DGNNs. Therefore, the choice of model must be carefully evaluated to determine whether the amount of data available justifies the use of more sophisticated models, considering the performance and technical debt of running them.

Finally, the practical application of the model depends heavily on the additional parameters, both controllable and external factors, which can at best be estimated. We conclude that a proper analysis of the cost of the target campaign, together with the cost of referral marketing, should be conducted to find the optimal threshold for final deployment. In our case, we recommend a careful study of the minimum reward necessary to induce a probability of response, using the proposed methods to evaluate the results. A/B testing and similar strategies can be used to fine-tune the value of the incentive.

While this study provides valuable insights into the applicability of DGNNs for influencer prediction, it is important to acknowledge its limitations. The approach used to construct the networks can influence model performance. Investigating the impact of network topology on performance is left for future research. In addition, the profit-driven evaluation framework presented in this study relies on the probability of referring without being targeted, which requires explicit modeling. Future research can focus on modeling this probability.

APPENDIX A NODE FEATURES

TABLE 6. Node features description.

Feature group description	Type	#feat.
# no payments from the contract creation date to the payment date	Numerical	1
Outstanding balance of a credit card at a cut-off date	Numerical	1
Payments between:	Numerical	2
1) start date of the billing period and cut-off date (a final date of the billing period):		
2) cut-off date to payment date (the date when a customer is supposed to pay)		
Credit card consumption in delivery app (total, average and standard deviation) between:	Numerical	6
1) start date to cut-off date;		
2) cut-off date to payment date		
Credit card consumption in other commercials (total, average and standard deviation) between:	Numerical	6
1) start date to cut-off date;		
2) cut-off date to payment date		
Number of transactions between:	Numerical	2
1) start date to cut-off date;		
2) cut-off date to payment date		
Credit limit	Numerical	1
Number of no payments between the contract creation date and the payment date	Numerical	1
Bureau score of the user at an application date	Numerical	1
Estimated income at an application date	Numerical	1
Amount of increase in the credit limit	Numerical	1
User's operative state at the payment date	Categorical	1
Credit limit increase indicator	Categorical	1
Counter of past referrals	Numerical	1

APPENDIX B NETWORK CHARACTERISTICS




FIGURE 7. Network characteristics.

APPENDIX C HYPERPARAMETER CONFIGURATIONS

TABLE 7. Model configurations: City 1.

Model	#layers		#emb.		#hid.dim.	SMOTE	#layers		#hid.dim.	Loss
	GNN	RNN	GNN	RNN	GNN	rate	RNN	RNN		
GIN-LSTM	2		512		512	0	2		256	Focal
GIN-GRU	2		512		256	0	2		256	BCE
GAT-LSTM	4		256		512	0	1		256	Focal
GAT-GRU	4		256		512	0	1		256	BCE
Features+PR-LSTM	/	/	/	/	/	0	1		512	Focal
Features+PR-GRU	/	/	/	/	/	0	1		512	Focal
Features-LSTM	/	/	/	/	/	0.5	1		256	Focal
Features-GRU	/	/	/	/	/	0	1		512	Focal
Static GAT	4		512		256	0.5	/	/	/	BCE
Static GIN	2		256		512	0.5	/	/	/	Focal

TABLE 8. Model configurations: City 2.

Model	#layers GNN	#emb. GNN	#hid.dim. GNN	SMOTE rate	#layers RNN	#hid.dim. RNN	Loss
GIN-LSTM	2	512	256	0	1	256	Focal
GIN-GRU	2	512	512	0	2	256	Focal
GAT-LSTM	2	256	512	0	2	512	Focal
GAT-GRU	4	256	512	0	2	256	Focal
Features+PR-LSTM	/	/	/	0	1	512	Focal
Features+PR-GRU	/	/	/	0	2	256	Focal
Features-LSTM	/	/	/	0	1	512	Focal
Features-GRU	/	/	/	0.5	1	256	Focal
Static GAT	4	256	256	0	/	/	BCE
Static GIN	2	512	256	0	/	/	Focal

TABLE 9. Model configurations: City 3.

Model	#layers GNN	#emb. GNN	#hid.dim. GNN	SMOTE rate	#layers RNN	#hid.dim. RNN	Loss
GIN-LSTM	2	512	512	0	1	256	Focal
GIN-GRU	2	512	512	0	1	512	Focal
GAT-LSTM	2	256	256	0	1	512	Focal
GAT-GRU	2	256	256	0	2	256	Focal
Features+PR-LSTM	/	/	/	0	1	512	BCE
Features+PR-GRU	/	/	/	0.5	1	512	Focal
Features-LSTM	/	/	/	0.5	1	256	Focal
Features-GRU	/	/	/	0	1	512	BCE
Static GAT	4	512	256	0.5	/	/	Focal
Static GIN	2	512	512	0.5	/	/	BCE

ACKNOWLEDGMENT

The authors would like to thank the anonymous company that provided the data.

REFERENCES

- [1] T. Hussain and H. Shouno, "MAGRes-UNet: Improved medical image segmentation through a deep learning paradigm of multi-attention gated residual U-Net," *IEEE Access*, vol. 12, pp. 40290–40310, 2024.
- [2] M. Óskarsdóttir, W. Ahmed, K. Antonio, B. Baesens, R. Dendievel, T. Donas, and T. Reynkens, "Social network analytics for supervised fraud detection in insurance," *Risk Anal.*, vol. 42, no. 8, pp. 1872–1890, Aug. 2022.
- [3] W. Yang, J. Li, S. Tan, Y. Tan, and X. Lu, "A heterogeneous graph neural network model for list recommendation," *Knowl.-Based Syst.*, vol. 277, Oct. 2023, Art. no. 110822.
- [4] S. Sivasankari and G. Vadivu, "Tracing the fake news propagation path using social network analysis," *Soft Comput.*, vol. 26, no. 23, pp. 12883–12891, Dec. 2022.
- [5] S. Rani and M. Kumar, "Influential node detection and ranking with fusion of heterogeneous social media information," *IEEE Trans. Computat. Social Syst.*, vol. 10, no. 4, pp. 1852–1874, 2023, doi: 10.1109/TCSS.2022.3195525.
- [6] J.-Y. Puigbo, G. Sánchez-Hernández, M. Casabayó, and N. Agell, "Influencer detection approaches in social networks: A current state-of-the-art," in *Proc. CCIA*, 2014, pp. 261–264.
- [7] I. Roelens, "Advances in referral marketing using social networks," Ph.D. dissertation, Fac. Econ. Bus. Admin., Ghent Univ., Ghent, Belgium, 2018.
- [8] T. Vanderschueren, T. Verdonck, B. Baesens, and W. Verbeke, "Predict-then-optimize or predict-and-optimize? An empirical evaluation of cost-sensitive learning strategies," *Inf. Sci.*, vol. 594, pp. 400–415, May 2022.
- [9] B. Baesens, V. Van Vlasselaer, and W. Verbeke, *Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques: A Guide to Data Science for Fraud Detection*. Hoboken, NJ, USA: Wiley, 2015.
- [10] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, Jan. 2020.
- [11] C. Zheng, Q. Zhang, S. Young, and W. Wang, "On-demand influencer discovery on social media," in *Proc. 29th ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2020, pp. 2337–2340.
- [12] Y. Shi, P. Quan, Y. Xiao, M. Lei, and L. Niu, "Graph influence network," *IEEE Trans. Cybern.*, vol. 53, no. 10, pp. 6146–6159, 2023, doi: 10.1109/TCYB.2022.3164474.
- [13] J. Skarding, B. Gabrys, and K. Musial, "Foundations and modeling of dynamic networks using dynamic graph neural networks: A survey," *IEEE Access*, vol. 9, pp. 79143–79168, 2021.
- [14] M. Óskarsdóttir, T. Van Calster, B. Baesens, W. Lemahieu, and J. Vanthienen, "Time series for early churn detection: Using similarity based classification for dynamic networks," *Expert Syst. Appl.*, vol. 106, pp. 55–65, Sep. 2018.
- [15] T. Zhao, X. Zhang, and S. Wang, "GraphSMOTE: Imbalanced node classification on graphs with graph neural networks," in *Proc. 14th ACM Int. Conf. Web Search Data Mining*, Mar. 2021, pp. 833–841.
- [16] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 4–24, Jan. 2021.
- [17] Y. Zhu, F. Lyu, C. Hu, X. Chen, and X. Liu, "Encoder–decoder architecture for supervised dynamic graph learning: A survey," 2022, *arXiv:2203.10480*.
- [18] Y. Seo, M. Defferrard, P. Vandergheynst, and X. Bresson, "Structured sequence modeling with graph convolutional recurrent networks," in *Neural Information Processing*, L. Cheng, A. C. S. Leung, and S. Ozawa, Eds., Siem Reap, Cambodia. Cham, Switzerland: Springer, 2018, pp. 362–373.
- [19] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. Schardl, and C. Leiserson, "EvolveGCN: Evolving graph convolutional networks for dynamic graphs," in *Proc. AAAI Conf. Artif. Intell.*, Apr. 2020, vol. 34, no. 4, pp. 5363–5370.
- [20] G. Ryu and L. Feick, "A penny for your thoughts: Referral reward programs and referral likelihood," *J. Marketing*, vol. 71, no. 1, pp. 84–94, Jan. 2007.
- [21] G. Armelini, C. Barrot, and J. U. Becker, "Referral programs, customer value, and the relevance of dyadic characteristics," *Int. J. Res. Marketing*, vol. 32, no. 4, pp. 449–452, Dec. 2015.
- [22] P. De Meo, M. Levene, and A. Provetti, "Branching processes reveal influential nodes in social networks," *Inf. Sci.*, vol. 644, Oct. 2023, Art. no. 119201.
- [23] S. A. Ríos, F. Aguilera, J. D. Nuñez-Gonzalez, and M. Graña, "Semantically enhanced network analysis for influencer identification in online social networks," *Neurocomputing*, vols. 326–327, pp. 71–81, Jan. 2019.
- [24] T. Huynh, H. Nguyen, I. Zelinka, D. Dinh, and X. H. Pham, "Detecting the influencer on social networks using passion point and measures of information propagation," *Sustainability*, vol. 12, no. 7, p. 3064, Apr. 2020.
- [25] M. Jain, A. Jaswani, A. Mehra, and L. Mudgal, "EDGly: Detection of influential nodes using game theory," *Multimedia Tools Appl.*, vol. 81, no. 2, pp. 1625–1647, Jan. 2022.
- [26] S. F. Seyfosadat and R. Ravanmehr, "Systematic literature review on identifying influencers in social networks," *Artif. Intell. Rev.*, vol. 56, no. S1, pp. 567–660, Oct. 2023.
- [27] E.-Y. Yu, Y.-P. Wang, Y. Fu, D.-B. Chen, and M. Xie, "Identifying critical nodes in complex networks via graph convolutional networks," *Knowl.-Based Syst.*, vol. 198, Jun. 2020, Art. no. 105893.
- [28] G. Zhao, P. Jia, A. Zhou, and B. Zhang, "InfGCN: Identifying influential nodes in complex networks with graph convolutional networks," *Neurocomputing*, vol. 414, pp. 18–26, Nov. 2020.
- [29] E. Tiukhova, E. Penalzoza, M. Óskarsdóttir, H. Garcia, A. C. Bahnsen, B. Baesens, M. Snoeck, and C. Bravo, "Influencer detection with dynamic graph neural networks," in *Proc. NeurIPS Temporal Graph Learn. Workshop*, 2022. [Online]. Available: https://openreview.net/forum?id=LJeE02IC_6
- [30] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. Int. Conf. Learn. Represent.*, 2018. [Online]. Available: <https://openreview.net/forum?id=rjXmpkCZ>
- [31] S. Brody, U. Alon, and E. Yahav, "How attentive are graph attention networks?" in *Proc. Int. Conf. Learn. Represent.*, 2022. [Online]. Available: <https://openreview.net/forum?id=F72ximsx7C1>
- [32] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *Proc. Int. Conf. Learn. Represent.*, 2019, doi: <https://openreview.net/forum?id=ryGs6iA5Km>.
- [33] A. Leman and B. Weisfeiler, "A reduction of a graph to a canonical form and an algebra arising during this reduction," *Nauchno-Tekhnicheskaya Informatsiya*, vol. 2, no. 9, pp. 12–16, 1968.
- [34] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. Pande, and J. Leskovec, "Strategies for pre-training graph neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2020. [Online]. Available: <https://openreview.net/forum?id=HJIWJJSFDH>
- [35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

- [36] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2014, pp. 1724–1734.
- [37] M. Óskarsdóttir, C. Bravo, W. Verbeke, C. Sarraute, B. Baesens, and J. Vanthienen, "Social network analytics for churn prediction in telco: Model building, evaluation and network architecture," *Expert Syst. Appl.*, vol. 85, pp. 204–220, Nov. 2017.
- [38] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Comput. Netw. ISDN Syst.*, vol. 30, nos. 1–7, pp. 107–117, Apr. 1998.
- [39] C. Van Den Bulte, E. Bayer, B. Skiera, and P. Schmitt, "How customer referral programs turn social capital into economic capital," *J. Marketing Res.*, vol. 55, no. 1, pp. 132–146, Feb. 2018.
- [40] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-balanced loss based on effective number of samples," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9268–9277.
- [41] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Represent. (ICLR)*, San Diego, CA, USA, Y. Bengio and Y. LeCun, Eds., May 2015.
- [43] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch geometric," in *Proc. ICLR Workshop Represent. Learn. Graphs Manifolds*, 2019. [Online]. Available: <https://arxiv.org/abs/1903.02428>
- [44] J. Davis and M. Goadrich, "The relationship between precision-recall and ROC curves," in *Proc. 23rd Int. Conf. Mach. Learn.*, 2006, pp. 233–240.
- [45] L. F. Wolff Anthony, B. Kanding, and R. Selvan, "Carbontracker: Tracking and predicting the carbon footprint of training deep learning models," Jul. 2020, *arXiv:2007.03051*.
- [46] A. Fernández, S. García, M. Galar, R. C. Prati, B. Krawczyk, and F. Herrera, *Learning From Imbalanced Data Sets*, vol. 10. Cham, Switzerland: Springer, 2018, doi: [10.1007/978-3-319-98074-4](https://doi.org/10.1007/978-3-319-98074-4).



ELENA TIUKHOVA received the Master of Science degree in business and information systems engineering, with a focus on computer science. She is currently pursuing the Ph.D. degree with the Research Center for Information Systems Engineering, KU Leuven. Her research interests include using machine learning, deep learning, and XAI to improve decision making in marketing, credit risk, and learning analytics.



EMILIANO PENALOZA received the master's degree in statistics from the University of Waterloo. He is currently pursuing the dual Ph.D. degree in computer science with Mila Quebec and Université de Montréal. His current research interests include robust and generalizable learning algorithms for networks and recommendation systems.



MARÍA ÓSKARSDÓTTIR received the Ph.D. degree in business analytics from the Faculty of Economics and Business, KU Leuven. She is currently an Associate Professor with the Department of Computer Science, Reykjavik University, and an Adjunct Professor with Western University. She has over 40 publications in high-impact journals and conferences. Her research interests include the intersection of network science and machine learning, looking at practical applications of data science for decision making in various domains. She serves as an Editor for *Machine Learning*.



BART BAESENS is currently a Professor of data science with KU Leuven and a Lecturer with the University of Southampton. He has co-authored more than 250 scientific articles and ten books. He received the OR Society's Goodeve Medal for Best JORS Paper in 2016 and the EURO 2014 and EURO 2017 Award for Best EJOR Paper. He was also named one of the World's top educators in Data Science by CDO magazine in 2021 and has educated tens of thousands of data scientists across the globe. He also has his own online learning BlueCourses platform.



MONIQUE SNOECK (Member, IEEE) is currently a Full Professor with the Research Center for Information Systems Engineering (LIRIS), KU Leuven, and a Visiting Professor with UNamur. She has (co-)authored over 130 peer-reviewed articles, half of which peer-reviewed journal articles. Her research interests include conceptual modeling, enterprise modeling, requirements engineering, model-driven engineering, and business process management and the teaching and learning of these topics. She is involved in numerous conferences, such as CAiSE, PoEM, ER, and EMMSAD.



CRISTIÁN BRAVO is currently an Associate Professor and the Canada Research Chair of the Banking and Insurance Analytics, University of Western Ontario, Canada, where he directs the Banking Analytics Laboratory. He also serves as an Editorial Board Member for *Applied Soft Computing* and the *Journal of Business Analytics*.

• • •