Predicting Temporal Metrics in Dynamic Systems Using Machine Learning

Eduardo S. O. Marques[1], Saulo Moraes Villela[1,2(✉)], Heder Soares Bernardino[1,2], and Alex Borges Vieira[1,2]

[1] Graduate Program in Computational Modeling, Federal University of Juiz de Fora, Juiz de Fora, Brazil
eduardo.santos@estudante.ufjf.br

[2] Department of Computer Science, Federal University of Juiz de Fora, Juiz de Fora, Brazil
{saulo.moraes,heder.bernardino,alex.borges}@ufjf.br

Abstract. Understanding and anticipating the behavior of dynamic systems is essential to support decision-making in complex scenarios. Time-Varying Graphs (TVGs) provide a natural representation for such systems, enabling the analysis of temporal interaction patterns. In this context, temporal centrality metrics, such as Cover Time and Time-constrained Coverage, allow the characterization of information propagation dynamics and the identification of relevant time instants. However, the computation of these metrics typically depends on the complete knowledge of future interactions, which restricts their applicability in real-time or predictive scenarios. To address this limitation, this work formulates the prediction of temporal centrality metrics as a time series forecasting problem, where past observations are used to estimate future values without requiring full knowledge of the network evolution. We evaluate four forecasting models, including Linear Regression, Elastic-Net, Multilayer Perceptron (MLP), and Temporal Convolutional Networks (TCN), using both real-world and synthetic datasets. The results show that linear models outperform neural approaches, achieving error reductions of up to 30% in terms of MAE and RMSE. These findings suggest that the temporal dynamics of the analyzed metrics can be effectively captured by low-complexity models, providing an interpretable and computationally efficient alternative for forecasting global properties in time-varying networks.

Keywords: Temporal centrality metrics · Forecasting models · Time-Varying Graphs

1 Introduction

Understanding how systems evolve over time—whether on the scale of hours, days, or years—is essential for supporting decision-making in complex scenarios such as epidemic spreading, communication networks, and human interaction

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026, LNCS 16768, pp. 141–156, 2027.
https://doi.org/10.1007/978-3-032-30488-9_9

systems. These systems are often characterized by emergent collective phenomena, in which global dynamics cannot be fully explained solely by the properties of individual components.

In the context of complex networks, analyzing such systems requires specialized tools, conceptual frameworks, and interpretations of quantitative results [1]. Complex systems are widely applied across domains such as biology, social sciences, and engineering, where different systems may share common structural and dynamical patterns despite their distinct nature [2].

Network metrics play a central role in this analysis, as they enable the characterization of structural and dynamical properties. In particular, metrics that incorporate the temporal dimension allow not only identifying relevant entities but also determining when interventions are most effective. For instance, in epidemiological scenarios, disease spreading depends not only on biological factors but also on the temporal organization of human contacts [3]. Therefore, understanding how the structure of contact networks evolves over time is essential to anticipate how such processes unfold.

Time-Varying Graphs (TVGs) provide a natural framework to model these systems by explicitly representing temporal interactions. Within this framework, temporal centrality metrics allow identifying not only the most relevant nodes but also critical moments in the network dynamics. In this work, we focus on two representative temporal centrality metrics: Cover Time (CT), which measures how quickly a node can reach all others, and Time-constrained Coverage (TCC), which quantifies the fraction of nodes reached within a limited time window. These metrics are particularly useful for capturing how quickly information or influence can propagate through the network [4].

However, the computation of these metrics depends on the availability of future interactions in the network. In practical scenarios, such information is inherently unavailable, making the direct application of these metrics in real-time or decision-support contexts challenging [5]. Most existing approaches address this limitation by focusing on predicting links, node states, or reconstructing the future topology of the network [6]. While effective, these strategies introduce additional complexity and often require modeling fine-grained network dynamics.

In this context, machine learning techniques—particularly time-series forecasting methods—offer a promising alternative by enabling the estimation of future values of these metrics based solely on past observations, without requiring explicit reconstruction of future network topology [7]. This perspective allows shifting the problem from predicting individual connections to forecasting aggregate structural properties.

In this work, we propose to model the temporal evolution of centrality metrics in TVGs as a time-series forecasting problem. Instead of reconstructing the network, our approach directly targets global structural metrics by predicting global metrics, such as CT and TCC, based solely on historical observations. We evaluate four forecasting models—Linear Regression, ElasticNet, Multilayer Perceptron (MLP), and Temporal Convolutional Networks (TCN)—using both

a real-world dataset of hospital contact interactions and synthetic temporal networks [4,8]. The main contributions of this work are:

- The formulation of temporal centrality estimation as a time-series forecasting problem, avoiding the need for future network reconstruction;
- A comparative evaluation of four forecasting models applied to temporal centrality metrics;
- An experimental analysis using both real-world and synthetic datasets, highlighting the relationship between data characteristics and model performance;
- Evidence that low-complexity linear models can outperform deep learning approaches in this context, achieving significant reductions in prediction errors;

Experimental results show that linear models consistently outperform the neural architectures, achieving reductions of up to 30% in MAE and RMSE. These findings indicate that the temporal dynamics of the evaluated metrics can be effectively captured by low-complexity models, providing an interpretable and computationally efficient alternative for forecasting global properties of time-varying networks under partial information.

2 Background

This section introduces the fundamental concepts that underpin this study, covering the structure and dynamics of time-varying graphs, as well as temporal centrality metrics.

2.1 Time-Varying Graphs

A graph is defined as an ordered pair $G = (V, E)$, where V represents the set of vertices (nodes) and E the set of edges that connect pairs of vertices. The edges can be directed or undirected and may be weighted, representing costs, distances, or interaction intensities [9].

In many real systems, the interactions between elements do not remain constant over time. To capture this dynamic behavior, the Time-Varying Graph (TVG) model is used, in which the network topology can evolve over time. Intuitively, a TVG can be interpreted as an ordered sequence of static graphs, each representing a distinct point in time, also called a snapshot [5]. This representation allows for explicitly modeling the temporal order of interactions, a fundamental aspect for the analysis of dynamic processes in networks. Formally, a TVG can be represented as an object $H = (V, E, T)$, where V is the set of vertices, T is the set of time instants considered, and E is the set of edges that connect vertices associated with different time instants [10].

In this work, TVGs are modeled as Multi-Aspect Graphs (MAGs) [11], a general structure capable of representing multilayer, dynamic, and temporal networks. A MAG is defined as an object $H = (A, E)$, where A is a finite list of sets

called aspects, and E is the set of edges. In the specific case of a TVG, two main aspects are considered: the vertices and the time instants, that is, $|A| = 2$. Thus, the model can be represented as $H = (V, E, T)$, maintaining formal equivalence with the classical definition of TVGs.

From this modeling, the concept of a temporal node is defined as an ordered pair (u, t), where $u \in V(H)$ and $t \in T(H)$. The set of all temporal nodes is given by the Cartesian product $VT(H) = V(H) \times T(H)$. This representation allows one to demonstrate that a TVG modeled as a MAG is isomorphic to a static directed graph [12]. This result enables the direct application of classical metrics and algorithms for directed graphs to the analysis of structural properties and dynamic processes in time-varying networks.

2.2 Temporal Centrality Metrics

In time-varying networks, the significance of a dynamic process depends not only on where it begins, but also on when it initiates. In various real-world scenarios—such as the spread of epidemics, the diffusion of information in social networks, or the dissemination of failures in complex systems—the specific point in time at which a process is triggered can decisively influence its speed, reach, and impact [5]. In this context, temporal centrality metrics emerge as a natural extension of traditional centrality measures used in static graphs, shifting the focus from the importance of vertices to the relative importance of moments in time.

While classical centrality metrics produce a ranking of vertices according to different structural criteria—such as degree, closeness, or betweenness—temporal centrality metrics establish a ranking of time instants that reflects the potential of each moment to influence the evolution of dynamic processes in the network [13]. Different definitions of temporal centrality, therefore, capture distinct notions of time-associated importance and are suitable for various types of applications and analytical objectives.

Temporal metrics do not depend solely on the network's state at a single point in time but also on its evolution over time. Consider a TVG $H = (V, E, T)$, as defined in Sect. 2.1, where $T(H)$ represents the set of time instants in the graph. Let $N = |T(H)|$ be the total number of time points considered and $T(H) = \{t_0, t_1, \ldots, t_{N-1}\}$ the ordered temporal sequence of these instants. For each instant $t_i \in T(H)$, temporal centrality metrics assess the potential impact of processes initiated at that moment, taking into account not only the topology of the graph at t_i but also the connections established in subsequent instants [5].

The temporal centrality metrics adopted in this work were originally proposed by [4] and are particularly suitable for characterizing both the speed and reach of diffusion processes in time-varying graphs. While Cover Time emphasizes the number of temporal steps needed to reach a predefined fraction of the network, Time-constrained Coverage evaluates the fraction of nodes reached under a fixed temporal constraint. In both cases, the metric values depend not only on the network configuration at the initial instant t_i but also on its future evolution, reflecting the dynamic diffusion potential associated with different process start times. This temporal dependence allows the metric values to be

modeled as time series over the instants t_i, which are explored in the following sections for dataset construction and the evaluation of forecasting models.

The Cover Time metric (CT) measures the minimum number of time instants (or steps) required for a diffusion process to reach a given fraction of nodes in the network [4]. The metric $CT(t_i, \tau)$ returns the average number of time instants (steps) across all nodes for a diffusion, starting at instant t_i, to reach a certain fraction τ of the network, and is represented by:

$$CT(t_i, \tau) = \begin{cases} \frac{1}{|V(H))|} \sum_{u \in V(H)} d_t(t_i, u, \tau), \text{ if } \tau \text{ is reached,} \\ \infty, \qquad\qquad\qquad\qquad \text{otherwise,} \end{cases} \tag{1}$$

where $d_t(t_i, u, \tau)$ is the number of time instants (steps) required for a diffusion, starting at node u at instant t_i, to reach a fraction τ of the nodes. If the fraction τ of nodes is not reached for a diffusion process that starts at any node at instant t_i, then $CT(t_i, \tau)$ is considered to be ∞.

It is assumed that the diffusion process follows a breadth-first search; that is, a generic node (u, t_i) starts the diffusion by sending information to all its adjacent nodes. Then, the adjacent temporal nodes retransmit information to their own adjacent nodes in the next time instant, and so on. The information is spread until the fraction τ of nodes is reached, and the time required for diffusion is recorded. This evaluation is repeated for each node $(u, t_i) \in VT(H)$. The average diffusion time is the cover time centrality of the time instant t_i.

The Time-constrained Coverage metric (TCC) evaluates the network coverage achieved by a diffusion process after a limited number of time steps [4]. For a given time instant t_i, $TCC(t_i, \Phi)$ returns the average fraction of TVG nodes reached by the diffusion process in Φ time steps, and is represented by:

$$TCC(t_i, \Phi) = \frac{1}{|V(H))|^2} \sum_{u \in V(H)} d_t(t_i, u, \Phi) \tag{2}$$

where $d_c(t_i, u, \Phi)$ is the number of nodes reached from node u after Φ time steps, when the diffusion starts at instant t_i. For this problem, the same diffusion process defined for the CT metric is considered.

3 Related Work

The modeling of complex systems using TVGs has been widely investigated in the literature as a means of representing structural dynamics that cannot be adequately captured by static graphs. Foundational works have established the theoretical basis for the analysis of temporal networks, including formal definitions, structural properties, and time-aware metrics [5,10]. In this context, several classical complex network metrics have been extended to the temporal domain, emphasizing that the order, duration, and frequency of interactions play a fundamental role in network dynamics [14,15].

Several studies explore temporal networks to analyze patterns of contact, mobility, and the diffusion of processes, especially in scenarios involving human interaction and epidemiology [16]. However, most of these approaches focus on descriptive analyses or on simulating dynamic processes over the network, devoting little attention to the predictive modeling of the temporal evolution of structural metrics derived from these representations.

Surveys such as [17] and [18] provide comprehensive overviews of methods for dynamic graph representation learning, including recurrent, attention-based, and continuous-time models. Architectures such as Temporal Graph Networks (TGN) and Temporal Graph Attention Networks (TGAT) have been proposed to model complex temporal dependencies in evolving graphs, achieving strong performance in tasks like link prediction and node classification [19,20].

In the field of dynamic network estimation and construction, methods based on regularized logistic regression with temporal smoothing were developed to infer graphs that evolve over time [21]. Later works extended this line of research to online and non-stationary scenarios, formulating the learning of TVGs as composite optimization problems that balance fidelity to observed data, structural sparsity, and temporal smoothness [22].

With the rise of machine learning techniques, particularly deep learning, a large body of work has focused on learning representations from dynamic graphs. Works have explored spatiotemporal forecasting in graph-structured data, where the goal is to predict signals associated with nodes over time. For example, models such as Dyn-GWN explicitly incorporate dynamic graph structures into time series forecasting frameworks [23]. Additionally, reinforcement learning approaches have been employed to identify optimal strategies for information diffusion in temporal networks, demonstrating the predictive potential of temporal and structural patterns [24].

Despite these advances, the majority of existing approaches focus on predicting either the graph topology (e.g., link prediction) or node-level signals. In contrast, the prediction of global structural metrics in temporal networks remains relatively underexplored. In particular, temporal centrality measures, which capture the importance of nodes considering the temporal dimension of interactions, have been widely studied from an analytical perspective but are rarely addressed in a predictive setting.

In this work, we address this gap by formulating the prediction of temporal centrality metrics as a supervised time series regression problem. Instead of reconstructing the network topology or relying on complex graph neural architectures, we model the temporal evolution of these metrics directly using lagged features. This approach enables capturing temporal dependencies in a simple, interpretable, and computationally efficient manner, while providing competitive predictive performance. By focusing on global structural properties rather than local or topological reconstruction tasks, our work complements existing literature and offers a novel perspective on forecasting in time-varying graphs.

4 Materials and Methods

This section describes the materials and methods adopted in this study, including the datasets used, the modeling of time-varying graphs, the computation of temporal centrality metrics, the transformation of data into time series, and the forecasting models employed.

Figure 1 presents an overview of the proposed methodology. The process begins with the acquisition of raw temporal interaction data, followed by its modeling as TVGs. Temporal centrality metrics are then computed at each time instant, generating time series that represent the evolution of the network dynamics. These series are transformed into a supervised learning format using lagged features, enabling the application of forecasting models. Finally, the predictions are evaluated using linear and neural models.

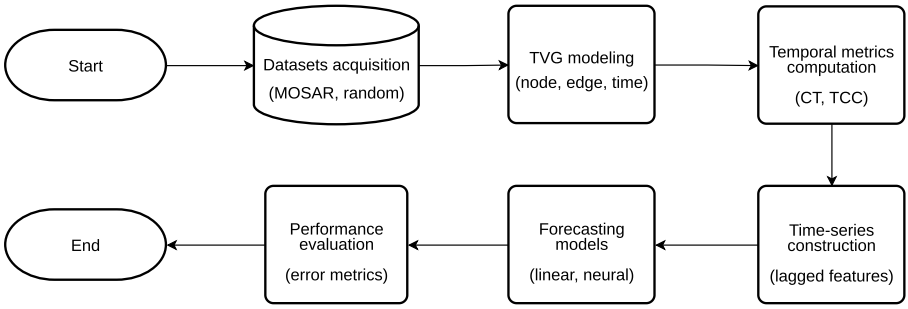

Fig. 1. Overview of the proposed methodology.

4.1 Datasets

The evaluation of temporal centrality metrics was carried out using a temporal graph (TVG) dataset collected in the context of the MOSAR project [8]. MOSAR (Mastering hOSpital Antimicrobial Resistance and its spread) is an international scientific collaboration initiative focused on analyzing the dynamics of transmission of antimicrobial-resistant bacteria in hospital environments. The dataset consists of records of face-to-face contacts between individuals—including healthcare professionals, staff, and patients—collected in a medical ward over two weeks, from July 25 to August 8, 2009.

In total, 160 volunteers participated in the study, each equipped with an RFID (Radio Frequency Identification) device capable of detecting other devices within an approximate radius of one meter. Contacts were recorded with a 30-second temporal granularity, enabling the inference of face-to-face interactions and the construction of network topology at each point in time.

The dataset was modeled as a TVG $M = (V, E, T)$, as defined in Sect. 2.1, where $V(M)$ represents the set of participants, $E(M)$ the face-to-face contacts,

and $T(M)$ the discretized time instants. In this context, $|V(M)| = 160$ and $|T(M)| = 40320$. The contacts were represented by undirected edges, modeled by including a directed edge and its reciprocal.

In addition to the real TVG, the metrics were also evaluated on a random TVG $R = (V, E, T)$, used as a synthetic reference for comparison. Each time instant corresponds to an independent Erdős–Rényi graph $G(n, p)$, maintaining the same number of vertices as the MOSAR TVG ($n = 160$) and considering $|T(R)| = 800$ time instants. The probability of edge creation was defined as $p = 0.01 \times \frac{\ln n}{n}$, ensuring a level of sparsity similar to that observed in the real network and mostly disconnected graphs, according to the classical connectivity threshold in random networks [25].

For both datasets, temporal metrics were calculated at each time instant, yielding time series that were later used in the forecasting stage. In this work, the results presented focus only on the finite values of the CT metric.

4.2 Time-Series Construction

The temporal centrality metrics were calculated at each discrete time instant in the network, yielding time series for the CT and TCC metrics. For this purpose, Python code was developed to generate the TVGs and calculate the corresponding metrics at each network snapshot. Next, only the time instants and metric values were extracted, stored in `.csv` files, and used as input for the machine learning models employed in the forecasting stage.

Preliminary experiments indicated that using time alone as the explanatory variable was insufficient to capture the dynamics of the analyzed metrics. In such cases, the models tended to learn only an approximately linear average trend, resulting in low predictive power. Given this limitation, the problem was reformulated as a time-series forecasting task using lagged features. This approach transforms the problem into a supervised format, in which past values of the series itself are used as explanatory variables for forecasting future values [26]. Furthermore, this strategy allows capturing nonlinear temporal relationships without imposing explicit assumptions of stationarity or linearity, and is widely adopted in machine learning models applied to time series [27].

Formally, let $y(t)$ be the result of a temporal centrality metric at a discrete instant t. For a given window size k, an attribute vector is constructed from the k immediately preceding values of the series. The prediction problem is then formulated as:

$$y(t+1) = f\big(y(t), y(t-1), \ldots, y(t-k+1)\big), \tag{3}$$

where $f(\cdot)$ represents a machine learning model trained to map the past values of the metric to its future value. The parameter k controls the amount of historical information provided to the model and directly influences its ability to capture short- and medium-term temporal dependencies. This formulation ensures that only information available up to instant t is used to predict $y(t+1)$, thus avoiding

any form of temporal data leakage, as recommended in best practices for time series predictive modeling [26].

Additionally, the separation between the training and test sets was carried out strictly in accordance with the temporal order of the data, so that the models were trained exclusively on past observations and evaluated on future instants, ensuring the validity of the evaluation process. By reformulating the problem in this way, it became possible to explicitly explore the temporal dependencies inherent to the temporal centrality metrics, resulting in significant gains in the predictive capacity of the models employed.

4.3 Datasets Parameters and Time Series Construction

The parameters used to construct the MOSAR and random datasets were defined based on previous studies to ensure methodological consistency and comparability of results [4]. For both networks, a fixed temporal window size of 30 time instants was adopted, which defines the resolution of the snapshots and the aggregation of events over time.

Not all available time instants were considered in the evaluation stage to ensure that the diffusion processes had a sufficiently wide temporal window to propagate throughout the TVGs, thereby avoiding truncation effects. For the MOSAR dataset, the parameters $\tau = 0.6$ and $\Phi = 2000$ were adopted, considering the first 33,000 time instants, although the network contains 40320 in total. Similarly, for the random dataset, the parameters $\tau = 0.6$ and $\Phi = 100$ were used, restricting the analysis to the first 660 instants out of a total of 800. The temporal centrality metrics were calculated at each time instant, yielding a time series.

To evaluate the time series using forecasting models, temporal windows of size k were used. The values were empirically set to $k = 15$ for the MOSAR and $k = 20$ for the random dataset, and were kept constant across both metrics in each network. Finally, the series was normalized to standardize the value scales and facilitate the learning process. The division between the training and test sets strictly preserved the temporal order of the data, with 80% of the initial observations used for training and the remaining 20% for testing.

4.4 Configuration and Evaluation of Prediction Models

The time series results were evaluated using different machine learning models for forecasting, including Linear Regression, Elastic Net, Multi-layer Perceptron (MLP), and Temporal Convolutional Network (TCN). All experiments were implemented in Python (version 3.12), using the Google Colab environment.

Additional models were also investigated during the exploratory phase, including ARIMA, SARIMAX, Long Short-Term Memory (LSTM) networks, Random Forest, Gradient Boosting, and Support Vector Machines (SVM). However, these approaches did not achieve satisfactory performance in the analyzed context, either due to limited predictive capability, instability in the results, or difficulty in adequately capturing the temporal dependencies of the studied

metrics. Therefore, the results presented in this work focus on the models that demonstrated the best performance and greater consistency.

The quantitative evaluation of model performance was carried out using metrics widely employed in the regression and time series forecasting literature, namely: coefficient of determination (R^2), adjusted coefficient of determination (R^2_{adj}), mean absolute error (MAE), mean squared error (MSE), root mean squared error (RMSE), and mean absolute percentage error (MAPE). These metrics enable performance analysis from different perspectives, accounting for both the magnitude and variability of errors.

Hyperparameter selection aimed to balance predictive capacity, complexity control, and stability of results, guided by best practices reported in the literature and preliminary experiments. Linear Regression was used as the reference model, with intercept and tolerance set to 10^{-6}. The Elastic Net model was employed as a regularized extension of linear regression, combining ℓ_1 and ℓ_2 penalties with parameters $\alpha = 10^{-4}$ and $l_1_ratio = 0.1$, aiming to reduce the risk of overfitting and handle possible feature correlations.

MLP was adopted to capture nonlinear relationships among temporal lags, with two fully connected hidden layers of 64 and 32 neurons, respectively, a ReLU activation function, and a linear output layer. Training was performed with the Adam optimizer, a learning rate of 10^{-3}, and an MSE loss function, and early stopping was used to mitigate overfitting. The TCN was used to explore local temporal patterns via one-dimensional causal convolutions, thereby preserving temporal causality. The architecture consists of two convolutional layers with 32 and 16 filters, respectively, each with a kernel size of 3, followed by a fully connected output layer. As with the MLP, training was performed with the Adam optimizer and the MSE loss function.

5 Results

This section presents the results obtained from applying machine learning models to forecast temporal centrality metrics. The analysis is conducted considering both the real MOSAR dataset and a synthetic random dataset.

5.1 MOSAR Dataset

Figure 2 presents the comparison between the actual values and the forecasts generated by the models for the CT metric, while Table 1 summarizes the quantitative performance. All models achieved high coefficients of determination ($R^2 > 0.9996$). The Linear model obtained the lowest MAE (1.3488) and MSE (14.0302), followed by the MLP and ElasticNet models. The TCN model presented higher error values compared to the other approaches.

For the TCC metric, the results are presented in Fig. 3 and Table 2. All models achieved high predictive performance, with R^2 values above 0.9997. The Linear model achieved the lowest error values, followed closely by the TCN model. The MLP and ElasticNet models presented slightly higher errors.

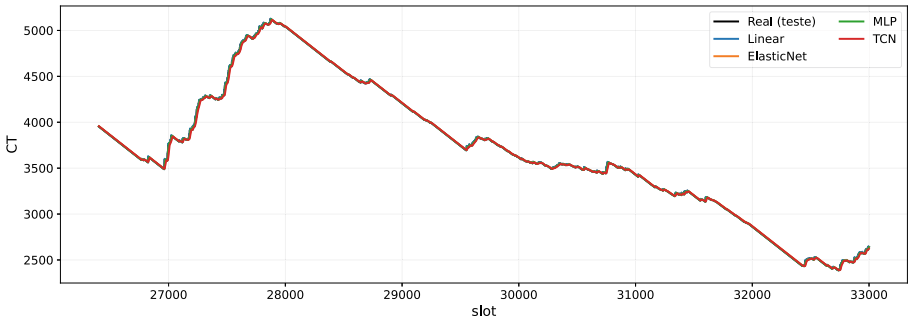

Fig. 2. Predictions from the models for the CT metric of the MOSAR dataset.

Table 1. Simulation results for the CT metric of the MOSAR dataset.

Model	R^2	MAE	MSE	R^2_{adj}	MAPE	RMSE
Linear	**0.999973**	**1.348768**	**14.030166**	**0.999973**	**0.000371**	**3.745686**
MLP	0.999917	3.698706	43.292463	0.999917	0.001015	6.579701
ElasticNet	0.999829	5.689566	89.036372	0.999829	0.001570	9.435909
TCN	0.999624	9.078452	195.793738	0.999623	0.002500	13.992632

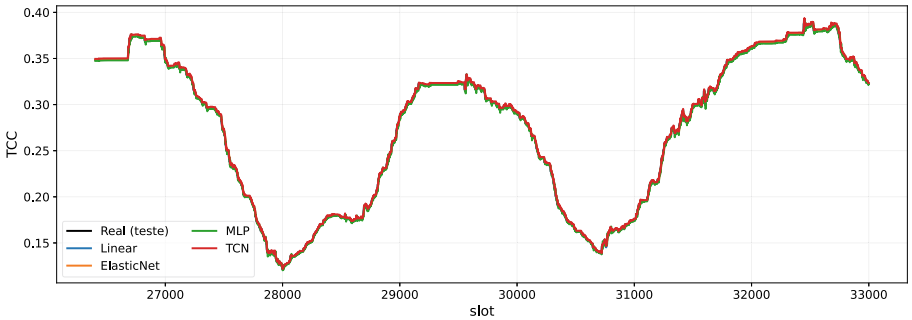

Fig. 3. Predictions from the models for the TCC metric of the MOSAR dataset.

Table 2. Simulation results for the TCC metric of the MOSAR dataset.

Model	R^2	MAE	MSE	R^2_{adj}	MAPE	RMSE
Linear	**0.999930**	**0.000270**	**4.568998e-07**	**0.999929**	**0.001150**	**0.000676**
MLP	0.999762	0.001073	1.544978e-06	0.999762	0.004080	0.001243
ElasticNet	0.999771	0.000659	1.489974e-06	0.999770	0.002782	0.001221
TCN	0.999869	0.000617	8.534101e-07	0.999868	0.002220	0.000924

5.2 Random Dataset

Figure 4 shows the comparison between actual and predicted values for the CT metric, while Table 3 presents the quantitative results. For the CT metric, the Linear and ElasticNet models achieved the highest R^2 values (0.979564 and 0.979236, respectively), along with the lowest error metrics. The MLP model also achieved competitive performance, although with slightly higher errors. The TCN model presented significantly lower performance, with $R^2 = 0.687641$ and higher error values.

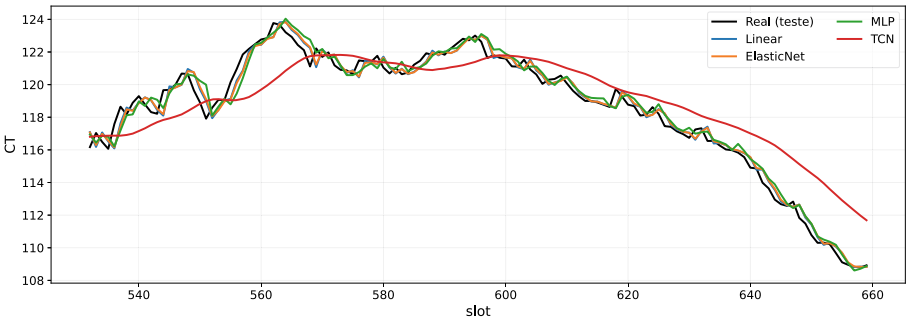

Fig. 4. Predictions from the models for the CT metric of the random dataset.

Table 3. Simulation results for the CT metric of the random dataset.

Model	R^2	MAE	MSE	R^2_{adj}	MAPE	RMSE
Linear	**0.979564**	**0.418950**	**0.284968**	**0.975745**	**0.003526**	**0.533824**
MLP	0.971709	0.493089	0.394507	0.966421	0.004158	0.628098
ElasticNet	0.979236	0.421180	0.289542	0.975355	0.003547	0.538091
TCN	0.687641	1.691775	4.355762	0.629256	0.014488	2.087046

For the TCC metric, the results are shown in Table 4 and Fig. 5. The ElasticNet model achieved the best performance, followed closely by the Linear model. The MLP and TCN models showed lower predictive accuracy, with higher error values across all evaluated metrics.

Table 4. Simulation results for the TCC metric of the random dataset.

Model	R^2	MAE	MSE	R^2_{adj}	MAPE	RMSE
Linear	0.946844	0.005697	0.000049	0.936908	0.014334	0.007029
MLP	0.886079	0.008413	0.000106	0.864785	0.021127	0.010290
ElasticNet	**0.949571**	**0.005524**	**0.000047**	**0.940145**	**0.013900**	**0.006847**
TCN	0.867562	0.009274	0.000123	0.842807	0.023401	0.011095

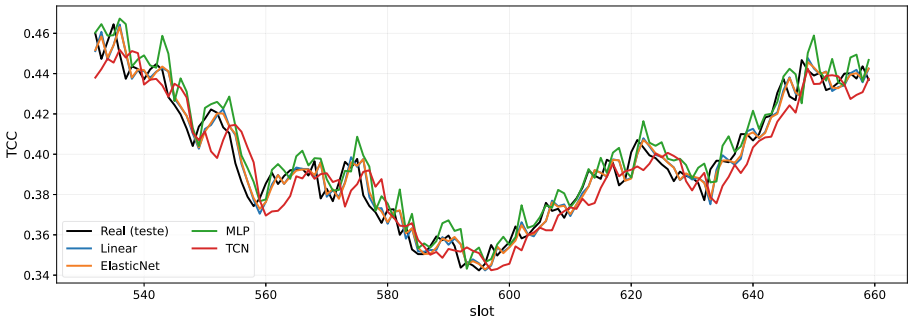

Fig. 5. Predictions from the models for the TCC metric of the random dataset.

6 Discussion

The results reveal distinct behaviors depending on the dataset characteristics and the evaluated temporal centrality metric. For the MOSAR dataset, all models achieved extremely high predictive performance, with R^2 values close to 1. This suggests that the temporal evolution of both CT and TCC metrics follows well-defined and stable patterns. In this context, linear models were sufficient to accurately capture these dynamics, outperforming more complex approaches such as MLP and TCN in terms of error minimization.

A possible explanation for this behavior lies in the structural properties of the MOSAR network, which may induce smoother and more predictable temporal relationships. As a consequence, increasing model complexity does not necessarily translate into performance gains, and may even introduce unnecessary variance.

In contrast, the random dataset exhibited a different behavior. Although linear models still achieved the best overall results, their performance was significantly lower compared to the MOSAR dataset, especially for the CT metric. This indicates that the temporal dynamics in random networks are less structured and more irregular, making them inherently more difficult to predict.

Furthermore, the poor performance of the TCN model in the random dataset suggests that deep temporal architectures are not universally suitable for this type of problem. This may be related to the absence of strong temporal dependencies or patterns that can be effectively exploited by convolutional structures.

Overall, the results highlight the importance of aligning model complexity with the underlying characteristics of the data. Simpler models not only achieved competitive or superior performance but also offer advantages in terms of interpretability and computational efficiency.

These findings have practical implications. In real-world applications where computational efficiency and interpretability are critical—such as monitoring dynamic systems or supporting decision-making—simpler models may provide a more reliable and cost-effective solution. This is particularly relevant when rapid predictions are required without access to future network information.

Despite these contributions, this study has some limitations. The analysis is restricted to two types of datasets and a specific set of temporal centrality metrics (CT and TCC). Additionally, the forecasting models were evaluated under a fixed experimental setup, which may not capture all possible configurations or parameter optimizations.

Future work may extend this approach by considering additional types of networks, including more heterogeneous and large-scale systems, as well as exploring other temporal metrics and hybrid modeling strategies. Moreover, integrating domain-specific knowledge or combining statistical and deep learning methods may further improve predictive performance in more complex scenarios.

7 Conclusion

This work investigated the feasibility of forecasting temporal centrality metrics in TVGs using machine learning models, considering both a real-world dataset (MOSAR) and a synthetic random network. The analyzed metrics (Cover Time and Time-constrained Coverage) enabled the assessment of the predictability of global dynamic properties across different temporal network structures.

The main contribution of the study lies in the formulation of temporal centrality prediction as a supervised time series forecasting problem, where lagged features are used to model temporal dependencies without requiring explicit knowledge of future network evolution. This perspective shifts the focus from predicting network topology or local interactions to forecasting global structural properties, providing a simple, interpretable, and computationally efficient alternative to more complex graph-based approaches.

Experimental results demonstrated that linear and regularized models, particularly Linear Regression and ElasticNet, consistently outperformed more complex neural architectures, achieving reductions of up to 30% in MAE and RMSE. These findings suggest that the temporal dynamics of the evaluated metrics exhibit patterns that can be effectively captured by low-complexity models, even in networks with evolving structures.

Overall, the proposed approach enables the estimation of relevant temporal network properties using only past observations, which is particularly valuable in scenarios where future interactions are inherently unknown, such as real-time monitoring of contact networks. This capability can support decision-making processes by providing timely insights into the evolution of information diffusion dynamics.

The study reinforces the importance of matching model complexity to the intrinsic characteristics of the data, showing that simpler approaches may be preferable even in dynamic and structured contexts such as temporal networks. For future work, we intend to extend the evaluation to a broader range of real-world temporal networks, incorporate additional metrics, and investigate more advanced validation strategies, such as statistical significance testing and cross-dataset generalization. Furthermore, exploring hybrid approaches that combine time series forecasting with graph-based learning methods may provide additional insights into the modeling of complex temporal systems.

Acknowledgments. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. The authors would also like to thank the Federal University of Juiz de Fora (UFJF) for the institutional and scientific support provided throughout this research.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Amaral, L.A., Ottino, J.M.: Complex networks: augmenting the framework for the study of complex systems. Eur. Phys. J. B **38**(2), 147–162 (2004)
2. Ma'ayan, A.: Complex systems biology. J. R. Soc. Interface **14**(134), 20170391 (2017)
3. Salathé, M., Freifeld, C.C., Mekaru, S.R., Tomasulo, A.F., Brownstein, J.S.: Influenza a (h7n9) and the importance of digital epidemiology. N. Engl. J. Med. **369**(5), 401 (2013)
4. Costa, E.C., Vieira, A.B., Wehmuth, K., Ziviani, A., Da Silva, A.P.C.: Time centrality in dynamic complex networks. Adv. Complex Syst. **18**(07n08), 1550023 (2015)
5. Holme, P., Saramäki, J.: Temporal networks. Phys. Rep. **519**(3), 97–125 (2012)
6. Lü, L., Zhou, T.: Link prediction in complex networks: a survey. XXPhys. A **390**(6), 1150–1170 (2011)
7. Box, G.E., Jenkins, G.M., Reinsel, G.C., Ljung, G.M.: Time Series Analysis: Forecasting and Control. Wiley, Hoboken (2015)
8. Lucet, J.C., et al.: Electronic sensors for assessing interactions between healthcare workers and patients under airborne precautions. PLoS ONE **7**(5), e37893 (2012)
9. Newman, M., Barabási, A.L., Watts, D.J.: The Structure and Dynamics of Networks. Princeton University Press (2011)
10. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. IJPEDS **27**(5), 387–408 (2012)
11. Wehmuth, K., Fleury, É., Ziviani, A.: On multiaspect graphs. Theoret. Comput. Sci. **651**, 50–61 (2016)
12. Wehmuth, K., Ziviani, A., Fleury, E.: A unifying model for representing time-varying graphs. In: Proceedings of the 2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA), pp. 1–10 (2015)
13. Kim, H., Anderson, R.: Temporal node centrality in complex networks. Phys. Rev. E-Stat. Nonlinear, Soft Matter Phys. **85**(2), 026107 (2012)
14. Holme, P.: Modern temporal network theory: a colloquium. Eur. Phys. J. B **88**(9), 1–30 (2015)
15. Battiston, F., et al.: Networks beyond pairwise interactions: structure and dynamics. Phys. Rep. **874**, 1–92 (2020)
16. Isella, L., Stehlé, J., Barrat, A., Cattuto, C., Pinton, J.F., Van den Broeck, W.: What's in a crowd? analysis of face-to-face behavioral networks. J. Theor. Biol. **271**(1), 166–180 (2011)
17. Kazemi, S.M., et al.: Representation learning for dynamic graphs: a survey. J. Mach. Learn. Res. **21**(70), 1–73 (2020)
18. Khoshraftar, S., An, A.: A survey on graph representation learning methods. ACM Trans. Intell. Syst. Technol. **15**(1), 1–55 (2024)

19. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., Bronstein, M.: Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637 (2020)
20. Xu, D., Ruan, C., Korpeoglu, E., Kumar, S., Achan, K.: Inductive representation learning on temporal graphs. arXiv preprint arXiv:2002.07962 (2020)
21. Kolar, M., Song, L., Ahmed, A., Xing, E.P.: Estimating time-varying networks. Ann. Appl. Stat. 94–123 (2010)
22. Natali, A., Isufi, E., Coutino, M., Leus, G.: Learning time-varying graphs from online data. IEEE Open J. Signal Process. **3**, 212–228 (2022)
23. Ibrahim, S., Tell, M., Mazumder, R.: Dyn-gwn: time-series forecasting using time-varying graphs with applications to finance and traffic prediction. In: Proceedings of the Fourth ACM International Conference on AI in Finance, pp. 167–175 (2023)
24. Mendonça, M.R., Barreto, A.M., Ziviani, A.: Efficient information diffusion in time-varying graphs through deep reinforcement learning. World Wide Web **25**(6), 2535–2560 (2022)
25. Erdős, P., Rényi, A.: On the evolution of random graphs. Publ. Math. Inst. Hung. Acad. Sci. **5**, 17–61 (1960)
26. Hyndman, R.J., Athanasopoulos, G.: Forecasting: principles and practice. OTexts (2018)
27. Franklin, J.: The elements of statistical learning: data mining, inference and prediction. Math. Intell. **27**(2), 83–85 (2005)