Predicting the Impact of Parameter Adjustments on Cellular Networks

Yongqian Sun[†¶], Qingliang Zhang[†], Yu Luo[†], Mingjie Li[‡*], Shenglin Zhang[†‖], Xiaolei Hua[‡], Renkai Yu[‡]
Xinwen Fan[‡], Lin Zhu[‡], Junlan Feng[‡], Dan Pei[§]
[†]*Nankai University*, {sunyongqian, zhangsl}@nankai.edu.cn, {1120250463, 2111934}@mail.nankai.edu.cn
[‡]*China Mobile Research Institute*, {limingjie, huaxiaolei, yurenkai, fanxinwen, zhulinyj, fengjunlan}@chinamobile.com
[§]*Tsinghua University*, peidan@tsinghua.edu.cn
[¶]*Tianjin Key Laboratory of Software Experience and Human Computer Interaction*
[‖]*Haihe Laboratory of Information Technology Application Innovation*

***Abstract*—Cellular networks are critical infrastructure for user equipment to access the Internet. Given the spatio-temporal dynamics of user distribution and traffic demand, operators adjust parameters such as *transmission power (TP)* and *cell individual offset (CIO)* to enhance network stability and service quality, However, predicting the impact of such adjustments is challenging due to limited historical adjustment data and complex metric dependencies. We propose *PIPCell*, a two-phase predictive framework. In phase one, *PIPCell* uses a closed-form multiplier from *TP* and *CIO* domain knowledge to calibrate adjustment-free *Workload* predictions from pre-trained Transformers. In phase two, a causal graphical model organizes multiple pre-trained Transformers to capture inter-metric dependencies and propagate adjustment effects. Experiments on real-world dataset from *China Mobile* show that *PIPCell* outperforms the best baseline by up to 25.8% in RMSE and 59.0% in sMAPE, demonstrating *PIPCell*'s potential for proactive and data-efficient cellular network optimization.**

***Index Terms*—cellular network, parameter adjustment, time series forecasting, causal inference**

I. Introduction

With the ongoing deployment of 5G, cellular networks have become the fundamental infrastructure for user access to the Internet, extensively implemented in next-generation NodeBs (gNBs). Given the spatio-temporal dynamics of user distribution and traffic demand, ensuring network stability and service quality is critical to maintaining a satisfactory user experience. To optimize network performance, operators typically adjust two key parameters: *Transmission Power (TP)* and *Cell Individual Offset (CIO)* [1]. However, this process heavily relies on manual expertise and often requires multiple iterative adjustments to achieve the desired outcomes [2]. Since operators usually do not directly control gNBs, there exists a feedback delay on the order of hours following parameter changes, resulting in prolonged intervals between consecutive adjustments. This significantly increases the overall time cost of optimization and delays improvements in service quality, thereby negatively impacting user experience. Therefore, accurately predicting the impact of parameter adjustments prior to deployment can effectively reduce the number of required tuning iterations, leading to more efficient network optimization and improved user satisfaction.

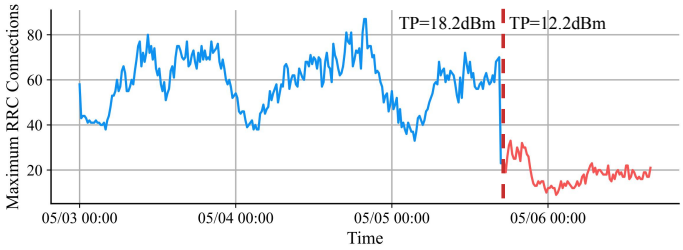

Fig. 1. Time series of the maximum number of RRC connections around an adjustment from a real-world cell. The dashed vertical line indicates the first data point with the new transmission power.

In this work, we develop a predictive framework to estimate the impact of parameter adjustments (*i.e.*, *TP* and *CIO*) on cell status metrics, with the goal of enabling operators to make informed decisions prior to real-world deployment. The cell status metrics, collected from a nationwide cellular network, are organized into three primary clusters encompassing 17 representative indicators, as summarized in Table I. Notably, *Workload* and *Quality of Service (QoS)* are the most critical for assessing overall network performance [3], [4]. *Interference* is out of a gNB's control and can also worsen *QoS*, measured as the noise level. Specifically, Fig. 1 illustrates an example of cell parameter adjustment, where the cell's maximum number of Radio Resource Control (RRC) connections—a key metric of *Workload*—experiences a significant decline after the *TP* is reduced by 6dBm. Accurately predicting such metric changes following parameter adjustments (*i.e.*, the right side of the vertical line in Fig. 1) is the interest of this work. Although prior work on traffic prediction [5], [6] and time-series forecasting [7] offers valuable insights into network trends, these methods typically assume static configurations and thus fail to capture the impact of dynamic cell parameter adjustments—an essential aspect of our scenario. As illustrated in Fig. 1, such adjustments can cause abrupt shifts in key metrics, making standard forecasting ineffective. Moreover, existing concept drift adaptation techniques [8], [9] mainly focus on updating models after a shift occurs, neglecting the challenge of predicting the impact of changes before they happen. These limitations highlight the need for a novel approach that directly

*Mingjie Li is the corresponding author.

TABLE I
CLUSTERS OF MONITORING METRICS

Cluster	Description	Monitoring Metrics
Workload	The number and behaviors of served UEs.	#(RRC connection established), #(E-RAB connection established), PRB downlink / uplink utilization, maximum #(RRC connection), average #(RRC connection), CCE utilization in PDCCH
Interference	The strength of irrelevant electromagnetic waves.	average noise level of the uplink PRBs
QoS	Quality of Service measurements.	drop rate, connection success rate, RRC connection success rate, E-RAB connection success rate, E-RAB drop rate (QCI=1), E-RAB connection success rate (QCI=1), handover success rate, VoLTE handover success rate, paging congestion rate

* Abbreviations: RRC (Radio Resource Control), E-RAB (Evolved Radio Access Bearer), PRB (Physical Resource Block), CCE (Control Channel Element), PDCCH (Physical Downlink Control Channel), QCI (QoS Class Identifier), VoLTE (Voice over LTE).

models the impact of parameter adjustments prior to their deployment.

Designing such a predictive framework, however, presents two key challenges due to the unique characteristics of parameter adjustment scenarios.

1) **Challenge 1:** Insufficient data for parameter adjustment scenarios. Historical data on parameter changes is limited, hindering the learning of generalizable patterns for impact prediction. Even when considering only two common parameters, *TP* from 0 to 60 dBm with 0.1 dB steps and *CIO* from minus 24 to 24 dB with 31 options [10], this results in over 18,000 unique combinations. This vast space far exceeds available data coverage, not to mention the over 2,000 tunable parameters per site in 5G networks [1].
2) **Challenge 2:** Modeling the complex dependencies among multiple metrics under cell parameter adjustments poses substantial challenges. In real-world scenarios, these dependencies are often dynamic and nonlinear, and highly sensitive to the system context. Under such conditions, building accurate and generalizable models becomes particularly difficult, especially when capturing the joint behavior of multiple metrics.

To tackle the aforementioned challenges, we propose *PIPCell*, a two-phase predictive framework that models the effect of cell parameter adjustments on key metrics. In the first phase, we generate adjustment-free predictions for *Workload* and *Interference*, followed by calibrating the *Workload* prediction through an multiplier that incorporates the effects of parameter adjustments. In the second phase, we employ a graphical model to organize multiple Transformers for explicitly modeling inter-metric dependencies, facilitating the propagation of parameter adjustment effects across other metrics. The contributions of this work are summarized as follows:

1) To the best of our knowledge, we are the first to propose a predictive framework that models the impact of cell parameter adjustments on cellular network time series. *PIPCell* enables proactive network optimization by forecasting how changes to parameters, *i.e.*, *TP* and *CIO*, influence key metrics.
2) To tackle limited parameter adjustment data, *PIPCell* derives a closed-form multiplier based on *TP* and *CIO*, to calibrate adjustment-free predictions, reducing dependence on scarce samples. To capture complex dependencies among metrics, *PIPCell* employs multiple Transformers organized by a graphical model, here causal parent predictions are embedded into the inputs of each Transformer to enable effective propagation of parameter effects.
3) We conduct extensive experiments on real-world dataset from *China Mobile*. Results show that *PIPCell* achieves notable improvements over competitive baselines—up to 25.8% and 59.0% gains in RMSE and sMAPE, respectively.

II. RELATED WORKS

Cellular Traffic Prediction. Early approaches like ABSENCE [11] and TOIP [12] model traffic using historical distributions within fixed time windows, ignoring temporal dependencies across time steps. Later works focus on near-future traffic prediction to enable proactive optimization such as load balancing [13]. Graph-based models [14]–[16] have become popular for capturing spatial-temporal correlations. We refer readers to [5] for a comprehensive survey. However, these methods generally assume a stable system configuration and fail to account for the impact of parameter adjustments on cellular traffic.

Time Series Forecasting. Traditional models such as ARIMA [17] and Prophet [18] rely on decomposition or linear modeling of recent history. Recent Transformer-based models [19]–[21] have shown strong performance by learning long-range dependencies. Similarly, these methods often assume a stable data generation process without external interventions [22], which limits their applicability in dynamic environments involving configuration changes.

Causal Inference. Estimating causal effects from observational data is challenging, as a single unit cannot undergo multiple interventions [23]. Tree-based methods like Causal Forest [24] approximate randomized trials by sample partitioning, but require large datasets and generalize poorly to unseen cases. Deep learning approaches attempt to model interventions [25] or latent variables [26], [27], yet often fail to capture the true data generation process [27]. To mitigate this issue, *PIPCell* incorporates domain knowledge to alleviate the model's dependence on observational patterns alone.

III. Homogeneous Cell Modeling

A. Modeling Workload by Area

A gNB typically equips multiple antennas, each covering a service area referred to as a cell. *TP* is a fundamental parameter of a cell, with higher power resulting in stronger Reference Signal Received Power (RSRP). The 3GPP-introduced *CIO* adjusts the effective RSRP in handover decisions by offsetting RSRP values to achieve load balancing [10]. According to the 3GPP standard, a UE tends to select cell t if:

$$RSRP_t + CIO_{t\to s} > RSRP_s + CIO_{s\to t} + Hys \tag{1}$$

where Hys is a hysteresis parameter designed to reduce frequent handovers at cell edges. Consequently, *TP* and *CIO* jointly influence UEs handover behavior and define cell boundary dynamics.

Due to the non-uniform distribution of UEs and the unavailability of detailed trajectory data for privacy reasons, we propose a trajectory-free model with a trade-off in accuracy. We assume that the average usage density of UEs at the cell edge is consistent with that of the entire cell. Based on this, we approximate *Workload* changes using the ratio of service area before and after adjustment, *i.e.*, $\frac{\|\mathbf{S}'\|}{\|\mathbf{S}\|}$, enabling adjustment-free time series forecasting to infer the impact on other performance metrics.

B. Cell Boundary Determination

To compute the cell area, it is necessary to define its boundary. Due to the hysteresis parameter Hys in (1), the boundary becomes a band with width. To achieve a precise boundary definition, we introduce a virtual boundary where the number of UEs connected to other cells inside the band equals those connected to the target cell outside, thereby factoring out the influence of Hys. This virtual boundary satisfies (2).

$$RSRP_t + CIO_{t\to s} = RSRP_s + CIO_{s\to t} \tag{2}$$

Since neighboring cells collectively affect the boundary, we fix their parameters to be identical, allowing controlled analysis of the target cell's influence. When *TP* increases from τ to $\tau+\delta$ dBm and *CIO* changes from o to $o+\delta_o$ dB, the received RSRP at the UE becomes $P'_R = \beta P_R$, where $\beta = 10^{0.1(\delta+\delta_o)}$. Here, β denotes the *TP* ratio between the target cell and its neighbors. Since electromagnetic wave power decays quadratically with distance, the coordinates (x, y) of a UE on the virtual boundary between two neighboring omnidirectional cells satisfy (3) based on (2), where P denotes the initial *TP* and R the inter-site distance. When $\beta \neq 1$, the boundary corresponds to an arc on the Apollonius circle [28]. Fig. 2 illustrates the boundary between two omnidirectional cells under $\beta < 1$. As for sectored cells, we treat the border between two neighboring cells of the same gNB as a ray starting from the gNB.

$$\underbrace{\frac{P}{(R-x)^2+y^2}}_{RSRP_t}\cdot\underbrace{10^{0.1o}}_{CIO_{t\to s}} = \underbrace{\frac{10^{0.1\delta}\cdot P}{x^2+y^2}}_{RSRP_s}\cdot\underbrace{10^{0.1(o+\delta_o)}}_{CIO_{s\to t}} \tag{3}$$

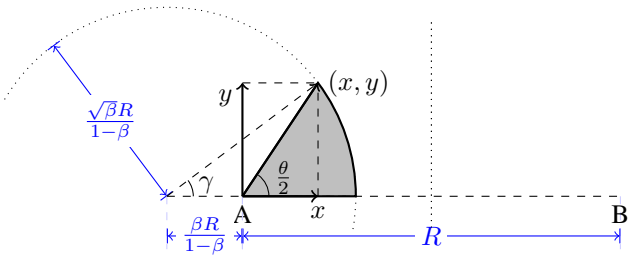

Fig. 2. Two omnidirectional cells. As A's *TP* changes from τ dBm to $\tau+\delta$ dBm and its *CIO* changes from o dB to $o+\delta_o$ dB, a UE connected to A will observe a change in the offset RSRP from P_R to $P'_R = \beta P_R$, where $\beta = 10^{0.1(\delta+\delta_o)}$. With $\beta < 1$, A's boundary shrinks to the dotted circle.

C. Parameters Change as a Multiplier

Let γ denote the central angle of the shaded circular arc, satisfying (4). The shaded area is then computed using (5):

$$\tan\frac{\theta}{2} = \frac{\frac{\sqrt{\beta}R}{1-\beta}\sin\gamma}{\frac{\sqrt{\beta}R}{1-\beta}\cos\gamma - \frac{\beta R}{1-\beta}} \Rightarrow \gamma = \frac{\theta}{2} - \arcsin\left(\sqrt{\beta}\sin\frac{\theta}{2}\right) \tag{4}$$

$$\begin{aligned} S_{shaded} &= \gamma\left(\frac{\sqrt{\beta}R}{1-\beta}\right)^2 - \frac{\sqrt{\beta}R}{1-\beta}\sin\gamma\cdot\frac{\beta R}{1-\beta} \\ &= \frac{\left(\gamma-\sqrt{\beta}\sin\gamma\right)\beta R^2}{(1-\beta)^2} \end{aligned} \tag{5}$$

The angle θ depends on gNB deployment, while β is derived from parameter adjustments. We calculate the ratio of areas after and before the adjustment: $\alpha(\beta \mid \theta) = S_{shaded}/\left(\frac{1}{4}R^2\tan\frac{\theta}{2}\right)$. For $\beta > 1$, we use symmetry: $\alpha(\beta \mid \theta) = 2 - \alpha\left(\frac{1}{\beta} \mid \theta\right)$. The complete expression is shown in (6), and we use $\alpha(\beta \mid \theta)$ to scale the predicted *Workload*: $\mathbf{W} \approx \tilde{\mathbf{W}}\alpha(\beta \mid \theta)$.

$$\alpha(\beta \mid \theta) = \begin{cases} \dfrac{4\beta\left(\gamma-\sqrt{\beta}\sin\gamma\right)}{(1-\beta)^2\tan\left(\frac{\theta}{2}\right)} & \text{if } \beta < 1 \\ 1 & \text{if } \beta = 1 \\ 2-\alpha\left(\frac{1}{\beta} \mid \theta\right) & \text{if } \beta > 1 \end{cases} \tag{6}$$

IV. PIPCell

A. Overview

We propose *PIPCell*, a framework built on a homogeneous cell model to capture the effects of parameter adjustments on multiple performance metrics in cellular networks. As shown in Fig. 3, *PIPCell* is first pretrained on adjustment-free data to learn fundamental time-series forecasting. During inference, it operates in two stages: adjusting and propagating. In the adjusting stage, *PIPCell* predicts adjustment-free *Workload* and *Interference*, and then refines the *Workload* prediction via the Homogeneous Cell Parameter Effect Predictor (*HCPEP*), alleviating the scarcity of adjustment data. In the propagating stage, *PIPCell* leverages Graphical Transformers (*GT*), where multiple Transformers are organized according to a dependency graph. Each Transformer conditions on parent predictions, enabling parameter-effect propagation across metrics and modeling the intricate dependencies introduced by parameter adjustments.

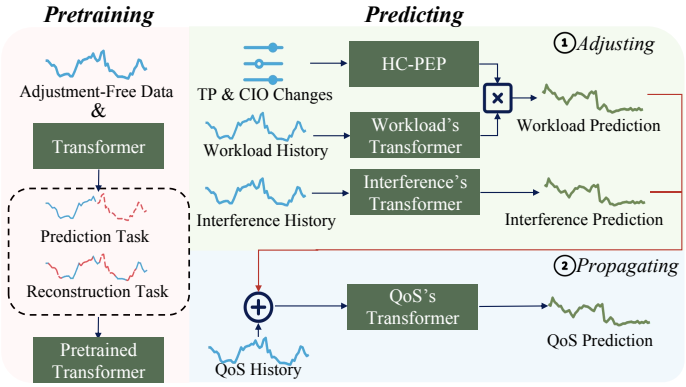

Fig. 3. Architecture of *PIPCell*.

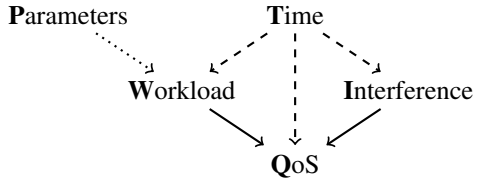

Fig. 4. Graphical model for a cell. Dashed arrows from time encode the periodicity of the related variables.

B. Pretraining

We adopt the Transformer architecture as the core prediction module, given its effectiveness in modeling long-range dependencies and complex temporal patterns in multivariate time series [19], [20], [29]. Multiple Transformers, each associated with a cluster of monitoring metrics, are pretrained on adjustment-free data to acquire fundamental forecasting capabilities. To enhance generalization, we leverage multi-task learning [30] by jointly optimizing time series forecasting and masked reconstruction. For forecasting, each Transformer predicts future values from historical inputs $\mathbf{x}^{t-I}, \cdots, \mathbf{x}^{t-1}$ with mean squared error (MSE) loss. For reconstruction, inspired by masked autoencoders [31], [32], we randomly mask parts of the input sequence and train the model to recover the masked values using MSE. The final objective sums both losses ((7)), where N, S, and d denote sample size, sequence length, and number of metrics, and $\omega_{i,j,k}$ indicates missing values (0 if missing, 1 otherwise).

$$loss = \frac{\sum_{i=1}^{N}\sum_{j=1}^{S}\sum_{k=1}^{d}(x_{i,j,k} - \tilde{x}_{i,j,k})^2\omega_{i,j,k}}{\sum_{i=1}^{N}\sum_{j=1}^{S}\sum_{k=1}^{d}\omega_{i,j,k}} \quad (7)$$

C. Adjusting

Based on Section III, we integrate the closed-form function $\alpha(\beta \mid \theta)$ into the Homogeneous Cell Parameter Effect Predictor (*HC-PEP*) module to explicitly model the impact of *TP* and *CIO* adjustments on *Workload*. As depicted in Fig. 3, following the Transformer's generation of adjustment-free *Workload* and *Interference* predictions, *HC-PEP* serves as a multiplier that scales the *Workload* prediction, thereby effectively incorporating the impact of parameter adjustments.

D. Propagating

We argue that parameter adjustments also indirectly affect metrics in other clusters beyond the *Workload* cluster, as shown in Fig. 4. To enable structured propagation of these effects, *PIPCell* employs a graphical model to organize multiple Transformers, referred to as Graphical Transformers (*GT*). After *HC-PEP* calibrates the *Workload* prediction to capture the direct effects of parameter changes, the metrics in each remaining cluster are predicted using a dedicated Transformer, which takes as input the historical data of the cluster and the prediction of its parent cluster—defined as the most causally relevant upstream component in the dependency graph. This design allows parameter effects to flow through the graph, influencing downstream metrics in a controlled and interpretable manner. Furthermore, the explicit graphical structure helps the *GT* avoid attending to spuriously correlated variables, thus improving learning efficiency. For example, *Interference* prediction is correctly modeled as independent of *Workload*, a constraint directly encoded in the graph.

V. Experiments

In this section, we evaluate the performance of *PIPCell* using a dataset collected by *China Mobile*, a nationwide cellular service provider. We aim to answer the following research questions (RQs):

RQ1: How does the performance of *PIPCell* compare to the baseline methods?

RQ2: Does each component of *PIPCell* contribute to its performance?

RQ3: How well does *PIPCell* generalize to cells in a different geographical region?

A. Experimental Setup

1) Datasets: We use a real-world dataset from *China Mobile* to evaluate different methods, denoted as $\mathcal{D}$. This dataset was collected from August 11th to 30th in 2022, including 1,045 times of cell parameter adjustments with distinct cells in the same geographical region. There are, on average, 628 data points per case before an adjustment. The sampling interval is 15 minutes. We select 17 monitoring metrics from 3 clusters shown in Table I.

2) Baseline methods and evaluation metrics: We employ a well-known machine learning model, GBRT [33], and two causal effect estimation methods, VCNet [25] and β-Intact-VAE [27], as baselines. The historical sequence serves as covariates, while the future time series represents causal effects. Prediction quality is evaluated for each monitoring metric individually using the original values prior to preprocessing, with root mean squared error (RMSE) and symmetric mean absolute percentage error (sMAPE) as evaluation metrics [5]. We randomly split the data into 50% for testing and 50% for training, repeating the process 10 times for robustness. The prediction horizon is one day (96 time steps), and we use the preceding 96 time steps as input to capture daily periodic patterns. Metrics are averaged over test cases within each trial, and the final results are reported as the mean across all trials.

TABLE II
PERFORMANCE EVALUATION OF PIPCELL AND BASELINE METHODS ON $\mathcal{D}$, AVERAGED AMONG MONITORING METRICS.

Method	Relative RMSE		Relative sMAPE	
	mean (std)	p-value	mean (std)	p-value
GBRT	1.347(0.80)	0.007*	2.440(2.69)	0.000*
VCNet	6.739(9.53)	0.000*	5.220(11.5)	0.000*
β-Intact-VAE	4.996(10.3)	0.002*	11.693(37.5)	0.000*
PIPCell	1.000(0.00)	/	1.000(0.00)	/

* statistically significant as $p < 0.05$

B. PIPCell vs. Baseline Methods (RQ1)

Table II presents a comprehensive comparison between *PIPCell* and three baseline methods: GBRT [33], VCNet [25], and β-Intact-VAE [27]. The results demonstrate that *PIPCell* consistently outperforms all baselines in terms of both RMSE and sMAPE, achieving improvements of 25.8% and 59.0%, respectively. Notably, although GBRT is a traditional machine learning approach, it surpasses the two deep learning-based methods, a phenomenon also observed in prior studies [7], [33]. *PIPCell* further advances performance beyond GBRT, primarily benefiting from the fine-grained modeling enabled by *GT* and *HC-PEP*.

To assess the statistical significance of the observed improvements, we conduct Wilcoxon signed-rank tests on the hypotheses that the relative RMSE (or sMAPE) is greater than 1, *i.e.*, that *PIPCell* achieves lower RMSE (or sMAPE) than a given baseline. The corresponding p-values are reported in Table II. Overall, the results confirm that *PIPCell* achieves statistically improvements over GBRT, VCNet, and β-Intact-VAE on both evaluation metrics, demonstrating its robustness and effectiveness across diverse monitoring scenarios.

C. Contribution of Key Components (RQ2)

We perform ablation studies to quantify the contributions of key components in *PIPCell*. As *HC-PEP* is specifically designed for *Workload* prediction, the graph-based model is essential for propagating the effect of parameter adjustments from the *Workload* to other metrics. Therefore, we keep the graph structure in all variants. Based on this design constraint, we focus our evaluation on two core components: (1) the *HC-PEP* multiplier and (2) the Transformer-based predictor. To this end, we construct five ablated variants of *PIPCell*: **A1**: Replace *HC-PEP* with the Multiplier module, which learns *Workload* adjustment ratios in a data-driven manner using GRU [34]-extracted historical features to model unobserved usage density. **A2**: Replace *HC-PEP* with VCNet. To assess the effect of the Transformer, we further substitute it with three architectures: **B1**: Replace the Transformer with GBRT. **B2**: Replace the Transformer with NLinear [7]. **B3**: Replace the Transformer with Autoformer [19].

1) Effect of HC-PEP: As shown in Table III, our method outperforms both the Multiplier-based (A1) and VCNet-based (A2) variants, demonstrating the effectiveness of *HC-PEP*. Compared to VCNet, *HC-PEP* improves RMSE and sMAPE by 38.3% and 50.1%, respectively. While Multiplier adopts a similar core idea as *HC-PEP*, it relies solely on data-driven learning. In contrast, *HC-PEP* integrates domain knowledge, achieving a lower sMAPE and competitive RMSE, thereby highlighting the benefit of domain-informed design.

TABLE III
PERFORMANCE EVALUATION OF PIPCELL'S VARIANTS ON $\mathcal{D}$, AVERAGED AMONG MONITORING METRICS.

Variant	Relative RMSE		Relative sMAPE	
	mean (std)	p-value	mean (std)	p-value
PIPCell	1.000(0.00)	/	1.000(0.00)	/
A1	0.973(0.26)	0.189	1.409(0.89)	0.001*
A2	1.621(0.99)	0.003*	2.005(1.34)	0.000*
B1	1.458(1.33)	0.013*	2.007(1.80)	0.000*
B2	5.219(13.7)	0.004*	13.400(44.9)	0.000*
B3	16.179(49.8)	0.001*	36.496(133.)	0.000*

* statistically significant as $p < 0.05$

TABLE IV
PERFORMANCE EVALUATION OF PIPCELL AND BASELINE METHODS IN ANOTHER REGION, AVERAGED AMONG MONITORING METRICS.

Method	Relative RMSE		Relative sMAPE	
	mean (std)	p-value	mean (std)	p-value
GBRT	1.665(0.76)	0.001*	3.294(7.88)	0.002*
VCNet	7.523(14.5)	0.000*	10.313(19.7)	0.000*
β-Intact-VAE	7.259(10.7)	0.000*	14.127(27.6)	0.000*
PIPCell	**1.000**(0.00)	/	**1.000**(0.00)	/

* statistically significant as $p < 0.05$

2) Effect of Transformer: To evaluate the effect of the Transformer in *PIPCell*, we compare it against three alternative variants: B1, B2, and B3. As shown in Table III, replacing the Transformer with GBRT (B1) and NLinear (B2) leads to performance degradation, demonstrating the superiority of the Transformer in time series forecasting within *PIPCell*. Although Autoformer [19] enhances the Transformer by incorporating time series decomposition, the variant using Autoformer (B3) still underperforms the original *PIPCell*. This can be attributed to the poor smoothness and weak periodicity inherent in cellular traffic data, as illustrated in Fig. 1. Moreover, for series that remain near zero for extended periods (*e.g.*, packet loss and congestion rates in $\mathcal{D}$), decomposition may introduce biased trend components, further impairing forecasting accuracy.

D. Generalization for Cells in Another Region (RQ3)

Due to variations in deployment scale and user density across regions, gNB performance metrics exhibit distinct temporal dynamics. To evaluate *PIPCell* 's cross-region generalization, we collected 60 test cases from a geographical area different from dataset $\mathcal{D}$. All models were trained solely on $\mathcal{D}$ and tested on these cases. As shown in Table IV, *PIPCell* consistently outperforms baselines in RMSE and sMAPE, demonstrating strong robustness and generalizability. We further illustrate a representative case on predicting the maximum number of RRC connections (Fig. 5). After *TP* decreased from 18.2 dBm to 12.2 dBm (vertical dashed line), the *GT* model, pretrained only on adjustment-free data, failed to capture the parameter change and persistently overestimated. In contrast, by applying the *HC-PEP* multiplier for

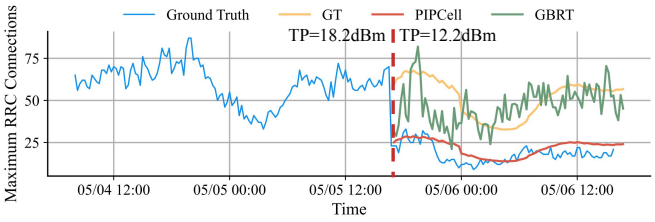

Fig. 5. The maximum number of RRC connections in a real-world case

dynamic correction, *PIPCell* adapts effectively to parameter adjustments and significantly improves accuracy. Compared with the strongest baseline (GBRT), *PIPCell* more faithfully tracks post-adjustment metric variations.

VI. Conclusion

We propose *PIPCell*, a two-phase framework for predicting the impact of cell parameter changes on time series of key performance metrics. In the first phase, *PIPCell* produces adjustment-free predictions of *Workload* and *Interference*, then refines *Workload* via the Homogeneous Cell Parameter Effect Predictor. In the second phase, a causal graphical model organizes multiple Transformers to capture inter-metric dependencies and propagate adjustment effects. Experiments on real-world datasets demonstrate that *PIPCell* reduces RMSE by up to 25.8% over the strongest baseline. Ablation studies highlight the contribution of each component, and cross-region evaluations confirm its generalization ability.

Acknowledgment

This work is supported by the National Natural Science Foundation of China (62272249, 62302244), and the Fundamental Research Funds for the Central Universities (XXX-63253249).

References

[1] J. Shodamola, U. Masood, M. Manalastas, and A. Imran, "A machine learning based framework for kpi maximization in emerging networks using mobility parameters," in *Proc. IEEE Int. Black Sea Conf. Commun. Netw. (BlackSeaCom)*, 2020.

[2] P. V. Klaine, M. A. Imran, O. Onireti, and R. D. Souza, "A survey of machine learning techniques applied to self-organizing cellular networks," *IEEE Commun. Surv. Tutor.*, vol. 19, no. 4, pp. 2392–2431, 2017.

[3] S. Kayyali, "Resource management and quality of service provisioning in 5g cellular networks," *arXiv preprint arXiv:2008.09601*, 2020.

[4] T. Zhou, Y. Huang, W. Huang, S. Li, Y. Sun, and L. Yang, "Qos-aware user association for load balancing in heterogeneous cellular networks," in *Proc. IEEE Veh. Technol. Conf. (VTC-Fall)*, 2014, pp. 1–5.

[5] W. Jiang, "Cellular traffic prediction with machine learning: A survey," *Expert Systems with Applications*, vol. 201, 2022.

[6] Y. Lin, Y. Gao, and W. Dong, "Bandwidth prediction for 5g cellular networks," in *Proc. IEEE/ACM Int. Symp. Quality of Service*, 2022.

[7] A. Zeng, M. Chen, L. Zhang, and Q. Xu, "Are transformers effective for time series forecasting?" in *Proc. AAAI Conf. Artif. Intell.*, 2023.

[8] J. a. Gama, I. Žliobaitundefined, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Comput. Surv.*, vol. 46, no. 4, mar 2014.

[9] M. Ma, S. Zhang, D. Pei, X. Huang, and H. Dai, "Robust and rapid adaptation for concept drift in software system anomaly detection," in *Proc. IEEE Int. Symp. Softw. Rel. Eng. (ISSRE)*, 2018, pp. 13–24.

[10] 3GPP, "NR; Radio Resource Control (RRC); Protocol specification," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 38.331, 2022, Release 17.

[11] B. Nguyen, Z. Ge, J. V. der Merwe, H. Yan, and J. Yates, "ABSENCE: Usage-Based Failure Detection in Mobile Networks," in *Proc. ACM Int. Conf. Mobile Comput. Netw. (MobiCom)*, 2015, pp. 464–476.

[12] S. Yang, Y. He, Z. Ge, D. Wang, and J. Xu, "Predictive impact analysis for designing a resilient cellular backhaul network," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 2, dec 2017.

[13] Q. Wang, C. He, K. Jaffrès-Runser, J. Huang, and Y. Xu, "Timely-throughput optimal scheduling for wireless flows with deep reinforcement learning," in *Proc. IEEE/ACM Int. Symp. Qual. Serv.*, 2022.

[14] X. Wang, Z. Zhou, Z. Yang, Y. Liu, and C. Peng, "Spatio-temporal analysis and prediction of cellular traffic in metropolis," in *Proc. IEEE Int. Conf. Netw. Protocols (ICNP)*, 2017.

[15] X. Wang, J. Zhao, L. Zhu, X. Zhou, Z. Li, J. Feng, C. Deng, and Y. Zhang, "Adaptive multi-receptive field spatial-temporal graph convolutional network for traffic forecasting," in *2021 IEEE Global Communications Conference (GLOBECOM)*, 2021.

[16] X. Zhou, Y. Zhang, Z. Li, X. Wang, J. Zhao, and Z. Zhang, "Large-scale cellular traffic prediction based on graph convolutional networks with transfer learning," *Neural Computing and Applications*, vol. 34, no. 7, pp. 5549–5559, jan 2022.

[17] X. Wang, Y. Kang, R. J. Hyndman, and F. Li, "Distributed arima models for ultra-long time series," *International Journal of Forecasting*, 2022.

[18] S. J. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, 2018.

[19] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," in *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 22 419–22 430.

[20] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," in *International Conference on Learning Representations*, 2023.

[21] D. Wu, X. Wang, Y. Qiao, Z. Wang, J. Jiang, S. Cui, and F. Wang, "Netllm: Adapting large language models for networking," in *Proceedings of the ACM SIGCOMM 2024 Conference*, 2024, p. 661–678.

[22] E. Bareinboim, J. D. Correa, D. Ibeling, and T. Icard, *On Pearl's Hierarchy and the Foundations of Causal Inference*, 1st ed. ACM, 2022, p. 507–556.

[23] W. Zhang, J. Li, and L. Liu, "A unified survey of treatment effect heterogeneity modelling and uplift modelling," *ACM Comput. Surv.*, vol. 54, no. 8, oct 2021.

[24] S. Wager and S. Athey, "Estimation and inference of heterogeneous treatment effects using random forests," *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1228–1242, 2018.

[25] L. Nie, M. Ye, qiang liu, and D. Nicolae, "VCNet and functional targeted regularization for learning causal effects of continuous treatments," in *International Conference on Learning Representations*, 2021.

[26] W. Zhang, L. Liu, and J. Li, "Treatment effect estimation with disentangled latent factors," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, may 2021, pp. 10 923–10 930.

[27] P. A. Wu and K. Fukumizu, "β-Intact-VAE: Identifying and estimating causal effects under limited overlap," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2022.

[28] H. Coxeter, "The problem of apollonius," *The American Mathematical Monthly*, vol. 75, no. 1, pp. 5–15, 1968.

[29] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting," in *Proc. 39th Int. Conf. Mach. Learn. (ICML)*, 2022, pp. 27 268–27 286.

[30] A. Navon, A. Shamsian, I. Achituve, H. Maron, K. Kawaguchi, G. Chechik, and E. Fetaya, "Multi-task learning as a bargaining game," in *Proc. 39th Int. Conf. Mach. Learn. (ICML)*, 2022, pp. 16 428–16 446.

[31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol. (NAACL-HLT)*, Jun. 2019, pp. 4171–4186.

[32] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 15 979–15 988.

[33] S. Elsayed, D. Thyssens, A. Rashed, H. S. Jomaa, and L. Schmidt-Thieme, "Do we really need deep learning models for time series forecasting?" 2021.

[34] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," 2014.