Modeling Coupled Transportation-Power Networks with Heterogeneous Graph Neural Networks

Tayab Uddin Wara, Saman Mazaheri Khamaneh, Tong Wu, George Atia
Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL, USA

***Abstract*—The integration of electric vehicles (EVs) into urban infrastructure creates unprecedented coupling between transportation and power distribution networks, where charging demand is simultaneously influenced by traffic patterns and grid constraints. Traditional forecasting approaches model these systems independently, failing to capture their bidirectional dependencies and resulting in inaccurate predictions. This work proposes a novel heterogeneous graph neural network (HGNN) framework for spatiotemporal forecasting in coupled power-transportation systems. We formulate a coupled optimization problem that captures realistic equilibrium dynamics between EV routing decisions and power dispatch constraints. Our HGNN architecture employs type-specific message passing and learned attention mechanisms to process heterogeneous features from road networks and power grids while preserving their structural coupling. Experimental results demonstrate that our framework significantly outperforms baseline approaches, reducing prediction error by over 40% compared to methods that model systems independently, validating the critical importance of capturing network heterogeneity and coupling dynamics for accurate charging demand forecasting.**

***Index Terms*—HGCN, coupled infrastructure networks**

I. Introduction

The rapid proliferation of electric vehicles (EVs) has fundamentally transformed urban infrastructure, with EV sales accounting for approximately 18% of global vehicle sales in 2024 and projected to exceed 30% by 2030 [1], [2]. This growth intensifies the coupling between transportation and power distribution networks, as EVs function as mobile loads that dynamically link traffic patterns with grid demand [3]. Unlike conventional vehicles, EV charging behavior is simultaneously influenced by transportation routing and power system constraints, creating bidirectional dependencies that traditional independent modeling fails to capture [4]. In Fig. 1, traffic congestion affects charging station accessibility and dwell times, while charging demand stresses distribution feeders and influences electricity pricing—which in turn shapes routing and charging decisions. These coupling effects become critical during peak hours when both networks operate near capacity, potentially triggering cascading failures [5].

Accurate prediction of charging station power consumption faces two fundamental challenges. First, the spatial correlation between road network topology and distribution feeder layouts creates interdependencies that cannot be captured by independent modeling [6]. Second, heterogeneous infrastructure components generate diverse measurements—from power system data (voltage, current, power flow) to transportation data (traffic flow, travel time, occupancy)—requiring unified processing frameworks [7]. Conventional forecasting methods, designed for homogeneous single-domain problems, cannot address these sorts of coupled, heterogeneous characteristics [7]. This motivates specialized machine learning approaches that leverage graph structures to capture both structural coupling and spatiotemporal dynamics.

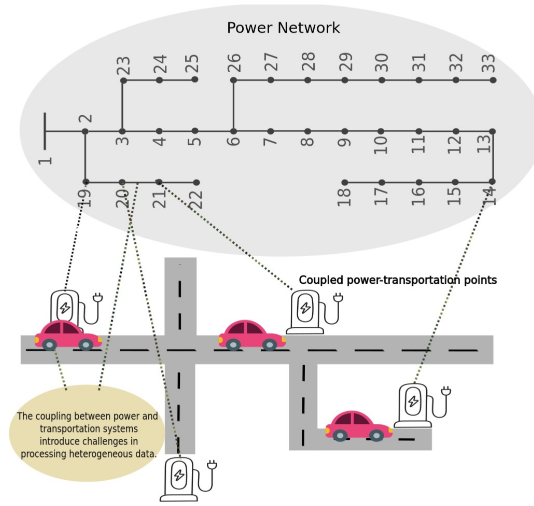

Fig. 1: Coupled power-transportation infrastructure for EV charging. The power distribution network (top) connects to the transportation network (bottom) through charging stations, creating bidirectional dependencies.

Existing approaches can be categorized into two streams. The first treats networks independently, where neural network methods for traffic prediction [8] or load forecasting [9] fail to capture bidirectional coupling effects. The second employs integrated optimization [10], but requires extensive domain knowledge, accurate parameters, and significant computational resources while often using simplified coupling representations that cannot capture complex spatiotemporal dynamics. Graph neural networks (GNNs) have emerged as data-driven alternatives for infrastructure modeling [11]. However, conventional homogeneous GNNs assume uniform node and edge types, making them unsuitable for systems with heterogeneous components and measurements [12]. While heterogeneous GNNs [13] enable type-specific transformations, their application to spatiotemporal forecasting in coupled networks remains unexplored. This gap motivates our specialized HGNN framework that addresses both structural heterogeneity and spatiotemporal dynamics for accurate charging station power consumption prediction. To address these challenges, this work makes the following contributions:

- We formulate a coupled transportation-power optimization problem modeling bidirectional dependencies between EV routing, charging decisions, and power constraints, generating a dataset that captures realistic coupling dynamics for data-driven modeling.

979-8-3315-8138-1

- We develop a novel HGNN architecture with type-specific transformations that process diverse measurements from heterogeneous components while preserving structural coupling, enabling accurate spatiotemporal forecasting of charging station power consumption.

The remainder of this paper is organized as follows: Section II introduces the problem formulation. Section III presents our HGNN framework. Section IV demonstrates effectiveness through case studies. Section V concludes the paper.

II. Problem Formulation

We address charging demand prediction in coupled transportation-power networks where EV routing depends on electricity prices while power dispatch depends on charging demand. Our approach generates coupled equilibrium solutions through iterative optimization, then trains an HGNN to predict demand without expensive re-solving.

A. Network Topology and Coupling Structure

Coupled infrastructure consists of transportation network $G_T = (V_T, E_T)$ and power network $G_P = (V_P, E_P)$ that couple at fast charging stations (FCS) $C \subseteq V_T \cap V_P$, where FCS $j \in C$ maps to both a transportation node and power bus. Time discretizes into hourly intervals $\mathcal{T} = \{1, 2, \ldots, T\}$. We formulate transportation routing and power dispatch models capturing these interdependencies through iterative coupling.

B. Transportation Network Model

Transportation demand originates from origin-destination (O-D) pairs $\mathcal{RS} := \{(r, s) : r, s \in V_T, r \neq s\}$, where q_t^{rs} (vehicles/hour) denotes demand at time t. EVs comprise fraction $\eta_t^{rs,E} \in [0, 1]$, with gasoline fraction defined as $\eta_t^{rs,G} := 1 - \eta_t^{rs,E}$ (stated upfront for mathematical rigor). Each O-D pair has non-charging paths $\mathcal{K}_n^{rs}$ (direct routes) and charging paths $\mathcal{K}_c^{rs}$ (routes with FCS stops), with complete path set $\mathcal{K}^{rs} := \mathcal{K}_n^{rs} \cup \mathcal{K}_c^{rs}$. For charging path $k \in \mathcal{K}_c^{rs}$, $\Omega_k^{rs} \subseteq C$ denotes accessible FCS locations. Path-link incidence $\delta_{k,a}^{rs} \in \{0, 1\}$ equals 1 if link $a \in E_T$ belongs to path k, else 0. Path flows $f_{k,t}^{rs,E}$ (EV vehicles/hour) and $f_{k,t}^{rs,G}$ (gasoline vehicles/hour) satisfy demand conservation via link flow aggregation:

$$x_{a,t} = \sum_{(r,s)\in\mathcal{RS}} \sum_{k\in\mathcal{K}^{rs}} (f_{k,t}^{rs,E} + f_{k,t}^{rs,G})\delta_{k,a}^{rs} \tag{1}$$

Travel time follows the Bureau of Public Roads (BPR) function:

$$t_{a,t}(x_{a,t}) = t_a^0 \left[1 + 0.15 \left(\frac{x_{a,t}}{c_a}\right)^4\right] \tag{2}$$

where t_a^0 (hours) is free-flow time and c_a (vehicles/hour) is capacity. EV battery feasibility requires initial state-of-charge (SOC) $S_0 \in [0, 1]$, consumption rate H_e (kWh/km), and range H_e^m (km). Distance l_a (km) denotes link length and l_{rj} (km) denotes distance from origin r to FCS j. Feasibility constraints depend on path type:

$$S_0 \geq \begin{cases} \sum_{a\in k} \frac{H_e l_a}{H_e^m} & k \in \mathcal{K}_n^{rs} \\ \min_{j\in\Omega_k^{rs}} \frac{H_e l_{rj}}{H_e^m} & k \in \mathcal{K}_c^{rs} \end{cases} \tag{3}$$

Travelers select routes via stochastic user equilibrium (SUE); generalized path cost combines travel time and expected charging delay:

$$C_{k,t}^{rs,g} = \omega \sum_{a\in k} t_{a,t}(x_{a,t})\delta_{k,a}^{rs} + \mathbf{1}_{g=E} \sum_{j\in\Omega_k^{rs}} \omega \frac{P_{k,t}^j}{\mu_m} \tag{4}$$

where ω (USD/hour) is value-of-time, μ_m (vehicles/hour) is FCS service rate, and $\mathbf{1}_{g=E}$ is indicator (1 for EV, 0 for gasoline). $P_{k,t}^j$ denotes FCS selection probability (formally defined below). SUE path flows are computed via entropy-regularized optimization:

$$\min_{\mathbf{f}, f_{k,t}^{rs,g}\geq 0} \quad \sum_{t,r,s,k,g} \left[\frac{f_{k,t}^{rs,g}}{\theta} \ln \frac{f_{k,t}^{rs,g}}{q_t^{rs}\eta_t^{rs,g}} + C_{k,t}^{rs,g} f_{k,t}^{rs,g}\right] \tag{5a}$$

$$\text{s.t.} \quad \sum_{k\in\mathcal{K}^{rs}} f_{k,t}^{rs,g} = q_t^{rs}\eta_t^{rs,g} \tag{5b}$$

$$x_{a,t} = \sum_{r,s,k} (f_{k,t}^{rs,E} + f_{k,t}^{rs,G})\delta_{k,a}^{rs} \tag{5c}$$

where $\theta > 0$ is route choice dispersion parameter (units: utility). Small θ concentrates flows on lowest-cost paths (deterministic behavior); large θ permits dispersed choices (heterogeneous behavior). Optimal solution satisfies multinomial logit: $f_{k,t}^{rs,g*} \propto \exp(-\theta C_{k,t}^{rs,g})$.

Upon reaching accessible FCS, EVs select a specific station via multinomial logit choice. Utility of FCS j is:

$$V_{k,t}^j = \beta_1 d_{k,t}^{rj} + \beta_2 \frac{1}{\mu_m} + \beta_3 c_{j,t} \tag{6}$$

where $d_{k,t}^{rj}$ (km) is detour distance, $\frac{1}{\mu_m}$ (hours/vehicle) is wait time, $c_{j,t}$ (USD/MWh) is electricity price, and $\beta_1, \beta_2, \beta_3 < 0$ are elasticities. Negative signs enforce rational consumer behavior: higher distance, wait time, and price each reduce utility. FCS selection probability is:

$$P_{k,t}^j = \frac{\exp(V_{k,t}^j)}{\sum_{j'\in\Omega_k^{rs}} \exp(V_{k,t}^{j'})} \tag{7}$$

defined only for accessible FCS $j \in \Omega_k^{rs}$.

Total FCS demand aggregates EV arrivals across all routes:

$$D_{j,t} = \sum_{(r,s)\in\mathcal{RS}} \sum_{\substack{k\in\mathcal{K}_c^{rs}: \\ j\in\Omega_k^{rs}}} f_{k,t}^{rs,E*} P_{k,t}^j \tag{8}$$

where $f_{k,t}^{rs,E*}$ is equilibrium EV flow from optimization (5). This equation forms the interface between transportation and power networks, converting routing decisions into aggregate charging demand at each FCS.

C. Power Distribution Network Model

Charging demand from Eq. (8) couples transportation routing to grid dispatch. The power network consists of buses $\mathcal{I} := V_P$ with generator buses $\mathcal{I}^g \subseteq \mathcal{I}$, and lines $\mathcal{L} := E_P$. For line $l \in \mathcal{L}$, sending and receiving buses are i_l^s and i_l^r. Incidence matrices $A_{l,i}^s, A_{l,i}^r \in \{0, 1\}$ indicate line origination/termination at bus i. Each bus i has base load $P_{i,t}^d$ (MW)

979-8-3315-8138-1

and $Q^d_{i,t}$ (MVAr), generation $p^g_{i,t} \in [p^{g,\min}_i, p^{g,\max}_i]$ (MW) and $q^g_{i,t} \in [q^{g,\min}_i, q^{g,\max}_i]$ (MVAr) for $i \in \mathcal{I}^g$, generation cost c^g_i (USD/MWh), and squared voltage $v_{i,t} \in [V^{\min,2}, V^{\max,2}]$ (per-unit^2) where $V^{\min,2} = 0.9025$ and $V^{\max,2} = 1.1025$. Each line l has power flows $pf_{l,t} \in [PF^{\min}_l, PF^{\max}_l]$ (MW) and $qf_{l,t} \in [QF^{\min}_l, QF^{\max}_l]$ (MVAr). FCS demand $D_{j,t}$ from (8) converts to electrical load $p^{\text{EV}}_{j,t} = D_{j,t} \cdot p_v$ (MW), where p_v (MW/vehicle). Total load at FCS bus $j \in C$ is $P^{\text{load}}_{j,t} = P^d_{j,t} + \mathbf{1}_{j\in C} p^{\text{EV}}_{j,t}$. Economic dispatch minimizes generation cost:

$$\min_{\mathbf{u}} \quad \sum_{i\in\mathcal{I}^g, t\in\mathcal{T}} c^g_i p^g_{i,t} \tag{9a}$$

$$\text{s.t.} \quad \sum_{l\in\mathcal{L}} pf_{l,t}(A^r_{l,i} - A^s_{l,i}) = P^d_{i,t} + \mathbf{1}_{i\in C} p^{\text{EV}}_{i,t} - p^g_{i,t} \tag{9b}$$

$$\sum_{l\in\mathcal{L}} qf_{l,t}(A^r_{l,i} - A^s_{l,i}) = Q^d_{i,t} - q^g_{i,t} \tag{9c}$$

$$v_{i^r_l,t} - v_{i^s_l,t} = 2(r_l pf_{l,t} + x_l qf_{l,t}) \tag{9d}$$

$$p^{g,\min}_i \le p^g_{i,t} \le p^{g,\max}_i,\ q^{g,\min}_i \le q^g_{i,t} \le q^{g,\max}_i \tag{9e}$$

$$PF^{\min}_l \le pf_{l,t} \le PF^{\max}_l,\ QF^{\min}_l \le qf_{l,t} \le QF^{\max}_l \tag{9f}$$

$$V^{\min,2} \le v_{i,t} \le V^{\max,2} \tag{9g}$$

for all $i \in \mathcal{I}$, $t \in \mathcal{T}$, $l \in \mathcal{L}$, and $i \in \mathcal{I}^g$ where applicable. Here $\mathbf{u} := \{p^g_{i,t}, q^g_{i,t}, v_{i,t}, pf_{l,t}, qf_{l,t}\}$ collects all decision variables. Constraint (9b) enforces active power balance with $\mathbf{1}_{i\in C}$ ensuring EV loads appear only at FCS buses. Constraint (9c) enforces reactive power balance. Constraint (9d) implements linearized DistFlow equations. Constraints (9e)–(9g) enforce generation, line flow, and voltage bounds.

Solving (9) yields optimal dispatch $\mathbf{u}^*$ and dual variables $\chi^*_{j,t}$ (locational marginal prices (LMPs), USD/MWh) from (9b) at FCS buses. Electricity prices incorporate congestion:

$$c_{j,t} = \chi^*_{j,t}\left(1 + \rho \cdot \mathbf{1}_{P^{\text{load}}_{j,t} > \bar{p}_j}\right) \tag{10}$$

where $\rho > 0$ is penalty multiplier and $\bar{p}_j$ (MW) is FCS capacity. When load exceeds capacity, price increases by a factor $(1+\rho)$, influencing routing through β_3.

D. Coupled Equilibrium and Data Generation

The coupled system exhibits bidirectional feedback: transportation routing (5) responds to power prices via (10), while power dispatch (9) responds to charging demand via (8). We solve for coupled equilibrium $(\mathbf{f}^*, \mathbf{u}^*)$ through iterative optimization: (i) initialize prices $c^{(0)}_{j,t}$; (ii) solve (5) given current prices, obtaining flows $\mathbf{f}^{(n)}$ and demand $\{D^{(n)}_{j,t}\}$; (iii) solve (9) given current demand, obtaining $\mathbf{u}^{(n)}$ and updated prices via (10); (iv) repeat until convergence $\|\mathbf{f}^{(n+1)} - \mathbf{f}^{(n)}\|_\infty < \epsilon = 10^{-4}$ vehicles/hour. This generates coupled equilibrium solutions across varying scenarios, forming the training dataset for the HGNN framework that predicts FCS demand from network features for inference.

III. Methodology

We train a heterogeneous graph neural network (HGNN) to learn the coupled equilibrium mapping from $\mathbf{f}^* \to D_j \to c_j \to \mathbf{f}$ (Eqs. (8)–(10)). Domain-specific message passing (P: power, T: transportation) with learned attention at FCS nodes preserves physics while capturing bidirectional coupling.

A. Heterogeneous Graph Construction and Message Passing

Construct heterogeneous graph $G_H = (V_H, E_H, \mathcal{R})$ from equilibrium $(\mathbf{f}^*_t, \mathbf{u}^*_t)$ at hour t, where $V_H := V_P \cup V_T$, $E_H := E_P \cup E_T$, $\mathcal{R} := \{\text{P}, \text{T}\}$. Power buses extract voltage state $\mathbf{x}^{\text{P}}_i(t) = [V^*_{i,t}; \arctan(q^{g*}_{i,t}/p^{g*}_{i,t})]^\top \in \mathbb{R}^2$ from dispatch; transportation nodes use learned embeddings $\mathbf{x}^{\text{T}}_u \in \mathbb{R}^{d_h}$ (initialized randomly, refined during training) since flows aggregate at links rather than node level. Power line features encode normalized magnitude and direction via doubly stochastic normalization (DSN, column-wise norm to $[0,1]$): $\mathbf{e}^{\text{P}}_l(t) = [\text{DSN}(|pf^*_{l,t}|), \mathbf{1}_{pf^*_{l,t}\ge 0}, \text{DSN}(|qf^*_{l,t}|), \mathbf{1}_{qf^*_{l,t}\ge 0}]^\top \in \mathbb{R}^4$ (indicator $\mathbf{1}_{(\cdot)}$ preserves flow polarity). Transportation links aggregate EV flows: $\mathbf{e}^{\text{T}}_a(t) = \text{DSN}(\sum_{(r,s)\in\mathcal{RS}} \sum_{k\in\mathcal{K}^{rs}_c: a\in k} f^{rs,E*}_{k,t}) \in \mathbb{R}$.

Over L HGNN layers, domain-specific aggregation processes each network. Power and transportation nodes compute messages combining node embeddings with averaged edge features:

$$\mathbf{m}^{\text{P}(\ell)}_i = \text{ReLU}\left(\mathbf{W}^{\text{P}}_{\text{agg}}\left[\mathbf{h}^{\text{P}(\ell-1)}_i \,\middle\|\, \frac{1}{|\mathcal{N}(i)|}\sum_{l\in\delta(i)} \mathbf{e}^{\text{P}}_l\right]\right) \tag{11}$$

$$\mathbf{m}^{\text{T}(\ell)}_u = \text{ReLU}\left(\mathbf{W}^{\text{T}}_{\text{agg}}\left[\mathbf{h}^{\text{T}(\ell-1)}_u \,\middle\|\, \frac{1}{|\mathcal{N}(u)|}\sum_{a\in\delta(u)} \mathbf{e}^{\text{T}}_a\right]\right) \tag{12}$$

where $\mathbf{h}^{(\ell)}_i \in \mathbb{R}^{d_h}$ denotes node embedding at layer ℓ (initialized as $\mathbf{h}^{(0)}_i = \mathbf{W}_{\text{embed}}\mathbf{x}_i$), $\mathbf{W}^P_{\text{agg}} \in \mathbb{R}^{d_h\times(d_h+4)}$, and $\mathbf{W}^T_{\text{agg}} \in \mathbb{R}^{d_h\times(d_h+1)}$ are learnable aggregation matrices, $\mathcal{N}(i)$ and $\delta(i)$ denote neighbors and incident edges.

Internal nodes preserve domain structure; FCS nodes $j \in C$ reconcile both constraints via learned attention. Project messages: $\mathbf{z}^{\text{P}}_j = \mathbf{W}^{\text{P}}_{\text{proj}}\mathbf{m}^{\text{P}(\ell)}_j$, $\mathbf{z}^{\text{T}}_j = \mathbf{W}^{\text{T}}_{\text{proj}}\mathbf{m}^{\text{T}(\ell)}_j$ (projected features $\mathbf{z}^k_j \in \mathbb{R}^{d_h}$); combine prior embeddings: $\mathbf{h}^{\text{state}}_j = \mathbf{W}_{\text{comb}}[\mathbf{h}^{\text{P}(\ell-1)}_j \,\|\, \mathbf{h}^{\text{T}(\ell-1)}_j]$ (state fusion capturing both priors). Compute attention via separate MLPs with LeakyReLU:

$$\boldsymbol{\alpha}_j = \text{softmax}\Big(\text{LeakyReLU}(\mathbf{a}_{\text{P}}[\mathbf{h}^{\text{state}}_j \,\|\, \mathbf{z}^{\text{P}}_j]),\ \text{LeakyReLU}(\mathbf{a}_{\text{T}}[\mathbf{h}^{\text{state}}_j \,\|\, \mathbf{z}^{\text{T}}_j])\Big) \tag{13}$$

where $\boldsymbol{\alpha}_j = [\alpha^{\text{P}}_j, \alpha^{\text{T}}_j]^\top$ are normalized attention weights determining which domain dominates. Fuse: $\mathbf{m}^{\text{fused}(\ell)}_j = \alpha^{\text{P}}_j\mathbf{z}^{\text{P}}_j + \alpha^{\text{T}}_j\mathbf{z}^{\text{T}}_j$ (attention-weighted fusion allows learned domain prioritization at FCS). Residual updates: $\mathbf{h}^{\text{P}(\ell)}_i = \mathbf{h}^{\text{P}(\ell-1)}_i + \text{ReLU}(\mathbf{m}^{\text{P}(\ell)}_i)$ for $i \in V_P \setminus C$; $\mathbf{h}^{(\ell)}_j = \mathbf{h}^{\text{P}(\ell-1)}_j + \text{ReLU}(\mathbf{m}^{\text{fused}(\ell)}_j)$ for $j \in C$.

B. Temporal Forecasting and Training

After HGNN layers, average power and transport FCS representations to create unified state (averaging preserves both perspectives equally): $\mathbf{h}^{(\text{fcs})}_j(t) = (\mathbf{h}^{\text{P}(L)}_j(t) + \mathbf{h}^{\text{T}(L)}_j(t))/2$

where $\mathbf{h}_j^{(\text{fcs})} \in \mathbb{R}^{d_h}$. Process K-hour lookback window $\{\mathbf{h}_j^{(\text{fcs})}(\tau)\}_{\tau=t-K+1}^{t}$ (chosen to capture weekly and daily patterns) through bidirectional LSTM (BiLSTM: multiple stacked layers with dropout): $\mathbf{s}_j = \text{BiLSTM}(\cdot) \in \mathbb{R}^{d_s}$ (concatenated forward/backward representations).

The temporal encoding $\mathbf{s}_j$ is processed through multi-layer perceptrons with two prediction heads for short-term and long-term forecasting. Generate H-step-ahead predictions (short and long horizons are concatenated and balanced for operational planning):

$$\widehat{D}_{j,t+1:t+H} = \begin{bmatrix} \mathcal{D}_j^{(\text{S})}(\mathbf{s}_j) \\ \mathcal{D}_j^{(\text{L})}(\mathbf{s}_j) \end{bmatrix} \in \mathbb{R}^H \tag{14}$$

where $\mathcal{D}_j^{(k)}(\mathbf{s}_j) := \mathbf{W}_{\text{out},j}^{(k)}\, \sigma\left(\mathbf{W}_{\text{mid2},j}^{(k)}\, \sigma\left(\mathbf{W}_{\text{mid1},j}^{(k)} \mathbf{s}_j\right)\right)$ for $k \in \{S, L\}$, with $\mathbf{W}_{\text{mid1},j} \in \mathbb{R}^{d_{m1} \times d_s}$, $\mathbf{W}_{\text{mid2},j} \in \mathbb{R}^{d_{m2} \times d_{m1}}$, $\mathbf{W}_{\text{out},j} \in \mathbb{R}^{(H/2) \times d_{m2}}$ and dropout on intermediate layers.

The model is trained by minimizing mean squared error:

$$\mathcal{L}(\theta) = \frac{1}{|S_{\text{train}}| \cdot |C| \cdot H} \sum_{(t,j) \in S_{\text{train}} \times C} \left\| \widehat{D}_{j,t+1:t+H} - D^*_{j,t+1:t+H} \right\|_2^2 \tag{15}$$

where $D^*_{j,t+1:t+H}$ are equilibrium targets from coupled problem. Optimize with Adam optimizer, gradient clipping (max norm), and early stopping (patience-based). The trained model predicts charging demand without re-solving coupled optimization (Eqs. (5a)–(9a)), enabling efficient spatiotemporal forecasting.

Algorithm 1 HGNN-BiLSTM Training

Input: Equilibrium solutions $\{(\mathbf{f}_t^*, \mathbf{u}_t^*)\}_{t=1}^T$; $\alpha = 0.0005$
Split data: 75% train, 15% validation, 10% test
for epoch $= 1, \ldots, 30$ **do**
  for each batch (size 16) **do**
    for HGNN layer $\ell = 1, 2, 3$ **do**
      Compute messages (Eqs. (11)–(12))
      For each FCS $j \in C$: compute attention (Eq. (13)), fuse, update
    end for
    Aggregate FCS embeddings; BiLSTM encoding; predict via Eq. (14)
    Compute loss Eq. (15); update θ with gradient clipping
  end for
  Validate on $\mathcal{D}_{\text{val}}$; stop if 8 epochs no improvement
end for
Output: Best trained parameters θ^*

IV. Experimental Setup and Results

A. Testbed and Dataset Generation

Coupled infrastructure comprises Sioux Falls transportation network (24 nodes, 76 links) and IEEE 33-bus power system, with 12 FCS at coupling locations $\{1, 2, 4, 5, 10, 11, 13, 14, 15, 19, 20, 21\}$ (Fig. 2). One-year hourly data ($T = 8{,}760$) with 10% EV penetration generated via Pyomo 6.9.4 and IPOPT 3.14.19, solving coupled traffic assignment and optimal power flow iteratively until convergence ($\|\mathbf{f}^{(n+1)} - \mathbf{f}^{(n)}\|_\infty < 10^{-4}$). Extracted features: transportation links aggregate EV flows $f_{k,t}^{rs,E}$ (76 dimensions),

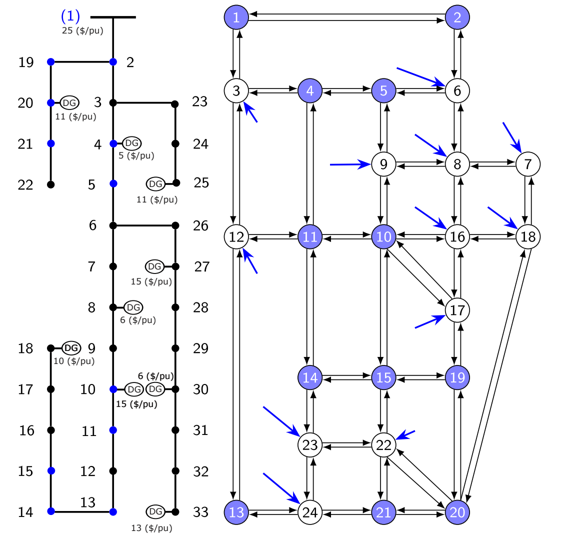

Fig. 2: Coupled power-transportation network.

FCS capture demand $D_{j,t}$ (12 dimensions), power buses encode voltage magnitude/angle (66 dimensions), transmission lines encode active/reactive flows (128 dimensions).

TABLE I: Performance Comparison

Model	MAE (MW)	RMSE (MW)
HGNN-LSTM (Ours)	**0.0651**	**0.0855**
Hom. GNN P	0.1222	0.1824
Hom. GNN T	0.1083	0.1825
FCNN	0.1193	0.1850
RNN	0.1275	0.1988
GRU	0.1005	0.1446
GAT	0.0994	0.1585
G. Transformer	0.1238	0.1789

B. Model Configuration and Baselines

HGNN-LSTM architecture: $L = 3$ heterogeneous layers ($d_h = 128$, aggregation matrices $\mathbb{R}^{128 \times 128}$), BiLSTM temporal module (2 stacked layers, 256 hidden units). Training: 75/15/10 split, $K = 72$-hour lookback, $H = 6$-hour horizon, Adam ($\alpha = 0.0005$), batch 16, gradient clipping (1.0), early stopping (patience 8, max 30 epochs). For baseline comparison, a homogeneous GNN using only power or transportation network features, matching the HGNN-LSTM capacity (128 hidden dimensions, 256 LSTM units, 3 GNN layers), is employed to isolate the benefit of heterogeneity. Additionally, FCNN, RNN, and GRU were used by flattening all 398 features (66 power node features + 256 power edge features + 76 transportation edge features) with no graph-like spatial awareness. For robust comparison, Graph Attention Network (GAT) with multi-head attention and homogeneous node processing was employed. Further, with the self-attention positional-encoder, Graph Transformer (G. Transformer) was implemented.

C. Evaluation and Results

Metrics across 12 FCS and 876 test hours ($N_t = 12 \times 876 \times 6$) with MAE, and RMSE. System-level performance

979-8-3315-8138-1

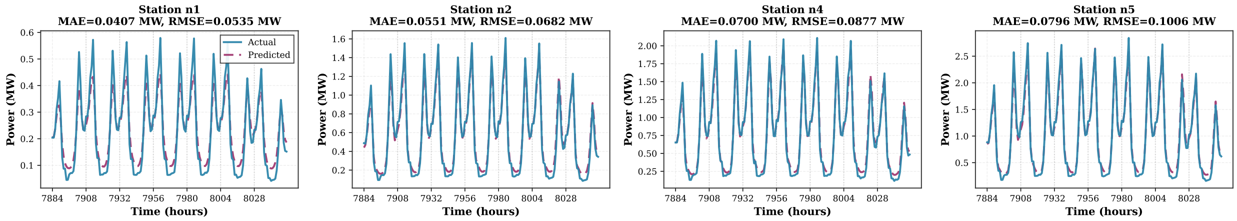

Fig. 3: One week actual vs predicted FCS power demand by predicted HGNN-LSTM (first 4 stations).

from Table I: HGNN-LSTM achieves MAE=0.0651 MW and RMSE=0.0855 MW, reducing error 46.7% and 53.1% versus homogeneous power-network baseline (Hom. GNN P: 0.1222 MW, 0.1824 MW) and 39.9% and 53.2% versus transportation-network baseline (Hom. GNN T: 0.1083 MW, 0.1825 MW). Indicating that both modalities are essential as heterogeneous fusion surpasses either alone. Sequential models lag 35–49% (GRU: 0.1005, FCNN: 0.1193, RNN: 0.1275), confirming spatial structure necessity. We even compared with advanced graph structures like GAT, achieving competitive results (MAE=0.0994 MW, RMSE=0.1585 MW), outperforming sequential baselines but falling 34.5% short of HGNN-LSTM, and G. Transformer with MAE and RMSE of 0.1238 MW and 0.1789 MW, respectively.Figure 3 demonstrates accurate daily peaks and weekly pattern reproduction for specific hours.

D. Online Deployment and Computational Efficiency

Rolling-horizon simulation in table II validates real-time applicability. HGNN-LSTM achieves 97.34 ms average inference per 6-hour horizon across 870 sequential predictions. Computational comparison yields 140× speedup vs. coupled optimization (84.69s vs. 11,814s), enabling the deployment plausibility of the model.

TABLE II: Online Deployment Performance

Inference Time		**Accuracy**	
Avg. (ms)	97.34	MAE (MW)	0.0679
Std. (ms)	2.28	RMSE (MW)	0.0902
99th %ile (ms)	102.44	Tests	870
Speedup: 140× vs. optimization (84.69s vs. 11,814s)			

TABLE III: Ablation Study: Component Contributions

Variant	**MAE (MW)**	**RMSE (MW)**
Full Model	**0.0651**	**0.0855**
w/o Attention	0.0715	0.0970
w/o Coupling	0.0818	0.1142
w/o Heterogeneity	0.0882	0.1167

E. Ablation Study

Component removal isolates individual contributions: (i) learned FCS attention, (ii) cross-domain coupling, and (iii) heterogeneous message passing. Table III quantifies these effects, revealing that heterogeneity removal causes the most severe degradation (+35.5% MAE), followed by coupling removal (+25.7%), then attention removal (+9.8%). This hierarchy establishes heterogeneous learning as the primary contribution, with cross-domain coupling and attention providing substantial complementary gains.

V. Conclusion

This work proposes a heterogeneous graph neural network framework for spatiotemporal forecasting of EV charging demand in coupled transportation-power networks. We formulate a coupled optimization problem capturing bidirectional dependencies between traffic routing and power dispatch, then develop an HGNN architecture with type-specific message passing and learned attention at charging station nodes. Combined with bidirectional LSTM for temporal modeling, our approach reduces prediction error by over 40% compared to independent network baselines. Future work will extend this framework to optimal control of EV charging and scheduling in coupled transportation-power systems.

References

[1] International Energy Agency, "Global ev outlook 2024," IEA, Tech. Rep., April 2024.

[2] Maybank Investment Banking Group, "Asean mobility: Poised for ev acceleration," Maybank, Tech. Rep., July 2023. [Online]. Available: https://www.maybank.com/iwov-resources/documents/pdf/blogs/2023/20230728-ASEAN-Mobility.pdf

[3] Y. Chen, S. Hu, Y. Zheng, S. Xie, Q. Hu, and Q. Yang, "Coordinated expansion planning of coupled power and transportation networks considering dynamic network equilibrium," *Appl. Energy*, vol. 360, p. 122789, 2024.

[4] S. G. Motlagh, J. Oladigbolu, and L. Li, "A review on electric vehicle charging station operation considering market dynamics and grid interaction," *Appl. Energy*, vol. 392, p. 126058, 2025.

[5] T. Wu, A. Scaglione, N. Ravi, S. Peisert, and D. Arnold, "Optimal control of differentially private ev charging: A scalable learning approach under uncertainty," *IEEE Internet Things J.*, 2025.

[6] S. Wang, C. Zhuge, C. Shao, P. Wang, X. Yang, and S. Wang, "Short-term electric vehicle charging demand prediction: A deep learning approach," *Appl. Energy*, vol. 340, p. 121032, 2023.

[7] J. Hu, X. Wang, and S. Tan, "Electric vehicle integration in coupled power distribution and transportation networks: A review," *Energies*, vol. 17, no. 19, p. 4775, 2024.

[8] D. A. Tedjopurnomo, Z. Bao, B. Zheng, F. M. Choudhury, and A. K. Qin, "A survey on modern deep neural network for traffic prediction: Trends, methods and challenges," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 4, pp. 1544–1561, 2020.

[9] H. Hou, C. Liu, Q. Wang, X. Wu, J. Tang, Y. Shi, and C. Xie, "Review of load forecasting based on artificial intelligence methodologies, models, and challenges," *Electr. Power Syst. Res.*, vol. 210, p. 108067, 2022.

[10] Z. Guo, F. Afifah, J. Qi, and S. Baghali, "A stochastic multiagent optimization framework for interdependent transportation and power system analyses," *IEEE Trans. Transp. Electrif.*, vol. 7, no. 3, pp. 1088–1098, 2021.

[11] T. Wu, A. Scaglione, A. P. Surani, D. Arnold, and S. Peisert, "Network-constrained reinforcement learning for optimal ev charging control," in *Proc. IEEE SmartGridComm*. IEEE, 2023, pp. 1–6.

[12] 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, 2021.

[13] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla, "Heterogeneous graph neural network," in *Proc. ACM SIGKDD*, 2019, pp. 793–803.

979-8-3315-8138-1