

A Reinforcement Learning-based Transmission Expansion Framework Considering Strategic Bidding in Electricity Markets

Tomonari Kanazawa, Hikaru Hoshino, Eiko Furutani

Department of Electrical Materials and Engineering,

University of Hyogo

Himeji, Japan

{er25q006@guh, hoshino@eng, furutani@eng}.u-hyogo.ac.jp

Abstract—Transmission expansion planning in electricity markets is tightly coupled with the strategic bidding behaviors of generation companies. This paper proposes a Reinforcement Learning (RL)-based co-optimization framework that simultaneously learns transmission investment decisions and generator bidding strategies within a unified training process. Based on a multi-agent RL framework for market simulation, the proposed method newly introduces a design policy layer that jointly optimizes continuous/discrete transmission expansion decisions together with strategic bidding policies. Through iterative interaction between market clearing and investment design, the framework effectively captures their mutual influence and achieves consistent co-optimization of expansion and bidding decisions. Case studies on the IEEE 30-bus system are provided for proof-of-concept validation of the proposed co-optimization framework.

Index Terms—Transmission Expansion, Electricity Markets, Reinforcement Learning, Multi-Agent Systems.

I. INTRODUCTION

The liberalization of electricity markets has led to increasing decentralization in investment and operation decisions among generators and transmission operators [1]. Generation is controlled by private Generation Companies (GENCOs), whose decisions are driven by their aim to maximize their profits. Transmission development and system (and market) operation are planned by independent entities, System Operators (SOs) and Market Operators (MOs). Under such competitive environments, transmission system operators are required to conduct *proactive planning* [2] that accounts for uncertainties in future generation and market outcomes. Traditional multi-level optimization models have addressed transmission expansion planning together with market operation (see, e.g., [3], [4]), yet these models typically assume rational market equilibrium and have limited flexibility to fully represent strategic bidding behaviors and price volatility.

Recent studies have applied deep Reinforcement Learning (RL) to planning of power systems in uncertain and decentralized environments. As reviewed in [5], existing studies have used RL in two distinct ways: operational planning and expansion planning. The first direction uses RL to simulate operation and market behavior, where agents representing

GENCOs learn strategic bidding policies that reproduce realistic price formation and market equilibria (see, e.g., [6]–[8]). These models capture short-term interactions among agents and enable analysis of market efficiency, price volatility, and regulatory impacts. The second direction uses RL or related learning techniques for expansion planning and investment, such as transmission expansion problems [9]–[11]. In these studies, RL serves as a high-level optimizer that explores design or investment decisions over long-term horizons, often neglecting operational details into simplified cost/reward functions. However, these two research lines have largely developed in parallel. Operational optimization or market-simulation approaches using RL provide realistic operational dynamics but typically assume fixed network configurations, while RL-based expansion planning approaches ignore how strategic bidding behaviors affect investment outcomes. In practice, transmission expansion planning is inherently coupled with market behavior: network upgrades influence locational prices and dispatch patterns, and conversely, bidding strategies and congestion patterns determine the economic value of transmission expansions.

To bridge this research gap, this paper proposes an RL-based co-optimization framework that incorporates strategic bidding and transmission expansion within a unified learning process. Based on the multi-agent deep RL framework for market simulations as reported in [8], we extend it to jointly optimize transmission investment decision as design variables that co-evolve with agent policies. This type of RL-based co-optimization of system design and its operational policy has been proposed primarily in the robotics community (see, e.g., [12], [13]) and recently applied to energy systems applications [14], [15]. To the best of our knowledge, this paper is the first to propose an RL-based co-optimization framework in the context of transmission expansion planning. Due to the well-known existence of multiple equilibria in joint transmission planning and market problems [3], [4], the proposed framework does not aim to identify a unique global optimum, but rather to obtain economically meaningful and locally optimal planner-market equilibria. Through case studies on the IEEE 30-bus system, we demonstrate that the proposed

This work was supported by research grant from Japan Power Academy.

framework captures the interaction between market dynamics and investment decisions, yielding cost-efficient expansion plans that are consistent with strategic market behavior.

II. PROBLEM FORMULATION

This study considers a two-level decision process that includes short-term market operation and long-term transmission investment. The formulation below defines the underlying market model and investment variables that form the basis of the proposed RL-based co-optimization framework.

A. Market Model

We consider a day-ahead electricity market consisting of N_g generating units (agents) participating through strategic bids. Each agent i aims to maximize its own profit by deciding a bid price $\lambda_i^{\text{bid}}(t)$ at each time step t based on observed states such as local demand and past market outcomes (the detailed modeling of bidding behavior is described in Sec. III-A). Given all submitted bids, the market operator performs DC Optimal Power Flow (DC-OPF) to clear the market under power balance and line-flow constraints at each time step t :

$$\min_{\{P_i, \theta_n\}} \sum_{i=1}^{N_g} \lambda_i^{\text{bid}}(t) P_i(t) \quad (1)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{G}_n} P_i(t) - D_n(t) = \sum_{(n,m) \in \mathcal{L}} B_{nm}(\theta_n(t) - \theta_m(t)), \quad \forall n \in \mathcal{N}_b, \quad (2)$$

$$|B_{nm}(\theta_n(t) - \theta_m(t))| \leq L_{nm}, \quad \forall (n,m) \in \mathcal{L}, \quad (3)$$

$$0 \leq P_i(t) \leq P_i^{\text{max}}, \quad \forall i \in \{1, \dots, N_g\}, \quad (4)$$

where $\mathcal{N}_b$, $\mathcal{L}$, and $\mathcal{G}_n$ are the sets of nodes (buses), lines, and generators connected to bus $n \in \mathcal{N}_b$, respectively; D_n is the demand at bus n ; θ_n is the voltage angle at bus n ; B_{nm} is the line susceptance; L_{nm} is the transmission capacity between buses n and m ; and P_i^{max} is the maximum generation capacity of unit i . Solving (1)–(4) yields the cleared generation $P_i^{\text{cleared}}(t)$ and nodal prices $\lambda_i^{\text{cleared}}(t)$, which reflect network congestion and marginal operating costs. The total operational system cost at time t is then defined as

$$C_{\text{oper}}(t) = \sum_{i=1}^{N_g} \lambda_i^{\text{cleared}}(t) P_i^{\text{cleared}}(t), \quad (5)$$

representing the total payment to generators in the cleared market. Although a single-block bidding structure is assumed here for simplicity, the formulation can be readily extended to multi-block or piecewise-linear bidding schemes, as commonly adopted in day-ahead markets.

B. Transmission Expansion

In this study, transmission expansion planning is considered for a single target year rather than a multi-stage horizon. That is, investment decisions are evaluated under representative operating conditions of a given year, and the resulting network configuration is assumed to remain fixed during the market

simulations. Let each transmission line $\ell \in \mathcal{L}$ have a base capacity L_ℓ^{base} , while in Sec. II-A the pair (n, m) refers to the sending and receiving buses of each line. We consider two types of formulations of a continuous expansion model that adjusts line capacities and a discrete siting model that determines whether to upgrade candidate lines:

$$\Delta L_\ell \in \mathbb{R}_+, \quad (\text{continuous capacity expansion}), \quad (6)$$

$$z_\ell \in \{0, 1\}, \quad (\text{discrete expansion decision}). \quad (7)$$

The effective line capacity is then expressed as

$$L_\ell = L_\ell^{\text{base}} + z_\ell \Delta L_\ell, \quad (8)$$

where either z_ℓ or ΔL_ℓ is fixed depending on the case study in Sec. IV. The investment cost for transmission expansion is modeled as:

$$C_{\text{exp}} = \sum_{\ell \in \mathcal{L}} c_\ell z_\ell \Delta L_\ell, \quad (9)$$

where c_ℓ is the expansion cost per unit capacity of the line ℓ per year. The overall optimization objective is to minimize the total cost including the expansion cost, defined as

$$J = W_{\text{anu}} \sum_{t=1}^T C_{\text{oper}}(t) + C_{\text{exp}}, \quad (10)$$

where W_{anu} stands for the annualization factor depending on the total time steps T . Such formulations are consistent with those employed in open-source capacity expansion models such as PyPSA [16] and other optimization frameworks.

III. RL-BASED CO-OPTIMIZATION FRAMEWORK

The overview of the proposed co-optimization framework is shown in Fig. 1. The upper part of the figure (highlighted in blue) represents the multi-agent RL environment used for market simulation and is described in Sec. III-A. The lower part illustrates the extension that enables simultaneous optimization of transmission capacity as explained in Sec. III-B.

A. Market Simulation Using Multi-agent RL

The competitive electricity market is modeled as a partially observable multi-agent environment largely based on [8]. Each GENCO, with a single generating unit, acts as an autonomous

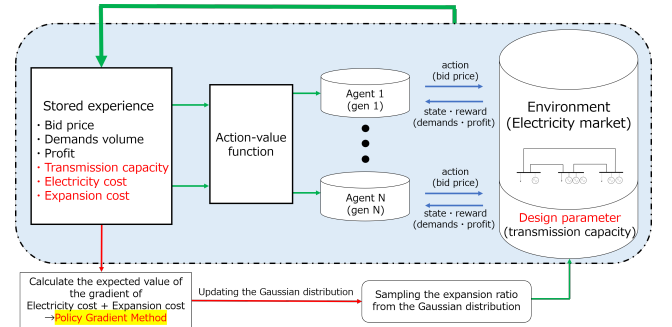


Fig. 1: Overview of the proposed co-optimization framework

agent submitting strategic bids to maximize its profit. At each time step t , agent i observes a local state given by

$$o_i(t) = [P_{\text{load}}(t), \lambda_i^{\text{bid}}(t-1), \omega^\top]^\top, \quad (11)$$

where $P_{\text{load}}(t) := \sum_{n \in \mathcal{N}_b} D_n(t)$ stands for the total system demand, and $\lambda_i^{\text{bid}}(t-1)$ is the bid submitted at the previous step. The vector ω represents the set of transmission design parameters that remain fixed during each episode, and its role in the co-optimization is detailed in Sec. III-B. Based on the observation $o_i(t)$, the agent i selects a normalized action $a_i(t) \in (0, 1)$, and the bid price $\lambda_i^{\text{bid}}(t)$ is calculated as

$$\lambda_i^{\text{bid}}(t) = \lambda_i^{\text{cost}}(t)(\alpha_i a_i(t) + 1) \quad (12)$$

where λ_i^{cost} is the marginal generation cost of unit i , and $\alpha_i > 0$ is a scaling parameter that defines the upper limit of λ_i^{bid} . When $a_i(t) = 0$, the agent bids truthfully at its marginal cost $\lambda_i^{\text{bid}}(t) = c_i(t)$, whereas larger $a_i(t)$ corresponds to higher bids. To avoid unrealistic temporal fluctuations in bidding, the bid prices are further constrained following [8] as

$$\frac{\max_t \lambda_i^{\text{bid}}(t)}{\min_t \lambda_i^{\text{bid}}(t)} \leq 1.5, \quad (13)$$

$$0.9 \leq \frac{\lambda_i^{\text{bid}}(t)}{\lambda_i^{\text{bid}}(t-1)} \leq 1.1. \quad (14)$$

After all agents submit their bids, the market operator clears the market as described in Sec. II-A, and each agent receives a reward r_i calculated as

$$r_i(t) = (\lambda_i^{\text{cleared}}(t) - \lambda_i^{\text{cost}}(t))P_i^{\text{cleared}}(t), \quad (15)$$

and updates its policy to maximize the long-term return

$$G_i = \sum_t \gamma^t r_i(t), \quad (16)$$

where $\gamma \in (0, 1)$ is the discount factor.

To address the coupled and non-stationary nature of multi-agent bidding, we employ the Multi-Agent Deep Deterministic Policy Gradient (MADDPG) algorithm. Each agent possesses an actor network $\pi_i(o_i; \theta_i^\pi)$ that maps local observations to continuous bid actions, and a centralized critic $Q_i(\mathbf{o}, \mathbf{a}; \theta_i^Q)$ that evaluates the joint action-value function using all agents' observations $\mathbf{o} = (o_1, \dots, o_{N_g})$ and actions $\mathbf{a} = (a_1, \dots, a_{N_g})$. During training, the centralized critic mitigates non-stationarity, while each actor learns its decentralized policy for execution. After each market clearing, experience tuples $(\mathbf{o}_t, \mathbf{a}_t, r_t, \mathbf{o}_{t+1})$ are stored in a replay buffer and used for off-policy updates of both actor and critic networks. This procedure allows the agents to learn Nash-like bidding equilibria consistent with network congestion and market clearing.

B. Co-optimization of Transmission Capacity

To enable co-optimization of transmission expansion, we treat the design vector $\omega = [\omega_1, \dots, \omega_{|\mathcal{L}|}]^\top$ as a random variable sampled at the beginning of each episode. Its sampling follows a design policy parametrized by $\mu = \{\mu_\ell\}_{\ell \in \mathcal{L}}$:

$$p_\mu(\omega) = \prod_{\ell \in \mathcal{L}} p_{\mu_\ell}(\omega_\ell), \quad (17)$$

Algorithm 1 Co-optimization of Bidding and Expansion

- 1: Initialize θ_i^π, θ_i^Q and parameter μ
 - 2: **for** episode = 1 to N **do**
 - 3: Sample design $\omega \sim p_\mu(\omega)$
 - 4: Simulate market episode with MADDPG bidding
 - 5: Compute episodic returns G_i
 - 6: Update actor and critic parameters θ_i^π, θ_i^Q
 - 7: **if** episode mod $N_{\text{up}} = 0$ **then**
 - 8: Update $\mu \leftarrow \mu + \alpha \nabla_\mu \mathbb{E}[G_{\text{total}}]$
 - 9: **end if**
 - 10: **end for**
-

where μ_ℓ is the learnable parameter governing the distribution of the design variable for line ℓ . In the continuous case, ω_ℓ corresponds to the capacity increment ΔL_ℓ , and we use a Gaussian distribution $\mathcal{N}(\mu_\ell, \sigma_\ell^2)$ as the design policy:

$$p_{\mu_\ell}(\omega_\ell) = \frac{1}{\sqrt{2\pi\sigma_\ell^2}} \exp \left[-\frac{(\omega_\ell - \mu_\ell)^2}{2\sigma_\ell^2} \right], \quad (18)$$

where σ_ℓ is a fixed standard deviation controlling exploration. In the discrete case, ω_ℓ corresponds to the binary line-upgrade decision $z_\ell \in \{0, 1\}$, for which we use a Bernoulli distribution:

$$p_{\mu_\ell}(\omega_\ell) = \mu_\ell \delta(\omega_\ell - 1) + (1 - \mu_\ell) \delta(\omega_\ell), \quad (19)$$

where the probability density is expressed using Dirac delta function $\delta(\cdot)$ for notational consistency, even though a discrete random variable is normally described by a probability mass function. The objective of the design-policy update is to minimize the total cost, and thus the episodic return G_{total} is defined as

$$G_{\text{total}} = -\left(W_{\text{anu}} \sum_{t=1}^T C_{\text{oper}}(t) + C_{\text{exp}}\right). \quad (20)$$

Based on this objective, the design-policy parameter μ is updated using the policy-gradient theorem [17]:

$$\nabla_\mu \mathbb{E}[G_{\text{total}}] = \mathbb{E} [\nabla_\mu \ln p_\mu(\omega) (G_{\text{total}} - \bar{G}_{\text{total}})], \quad (21)$$

where ∇_μ is the gradient operator with respect to μ , and $\bar{G}_{\text{total}}$ is a moving-average of G_{total} introduced to reduce the variance of the gradient estimator. The expectation is approximated by the sample average over the most recent N_{up} episode, and p_μ can be replaced with a probability mass function for discrete variables. The overall learning flow is summarized in Algorithm 1.

IV. NUMERICAL EXPERIMENTS

This section evaluates the proposed framework through proof-of-concept examples on the IEEE 30-bus system. After presenting the system settings in Sec. IV-A, two types of experiments are conducted. In Sec. IV-B, we compare the proposed method with a conventional two-stage benchmark using PyPSA [16] under continuous expansion for two selected lines. In Sec. IV-C, we examine a more realistic setup involving discrete siting decisions over multiple candidate lines.

A. System Settings

The topology of the IEEE 30-bus system is shown in Fig. 2. The generators at bus 2 (Gen 1), bus 23 (Gen 2), and bus 27 (Gen 3) are considered as strategic bidders that maximize their profits, and all other generators submit their true marginal cost. The marginal generation cost is assumed constant over the output range, and is set to \$50/MWh for strategic bidders and \$55/MWh for other generators. In this system, a relatively large demand is concentrated in the upper-right area, and the lines connecting to this region are selected as candidates for expansion. The baseline transmission capacities of lines 1-2, 3-4, 6-10, and 9-10 are set to 20 MW, while the capacities of lines 4-12 and 27-28 are set to 10 MW. The capacity expansion cost is assumed to be \$100,000/MW/year. The market simulations are performed with an hourly resolution over $T = 48$ intervals for illustration, though longer horizons are readily scalable within the RL framework.

The bidding agents are trained with MADDPG, where each actor network consists of six fully connected layers and each critic network consists of four layers. All hidden layers has 128 units and use ReLU activation. Adam learning rates are set to 10^{-7} (actor) and 10^{-5} (critic), with discount $\gamma = 0.99$, target-network smoothing $\tau = 5 \times 10^{-3}$, replay buffer size 2×10^4 , and minibatch size 64. Actions are bounded in $[0, 1]$, and the parameter α_i in (12) is set to 1. With this setting, the bidding policies of all strategic generators were observed to converge after approximately 1000–2000 episodes under a fixed transmission network configuration. Due to the coupled interaction between strategic bidding and transmission expansion, as well as stochastic exploration in RL training, bidding agents may converge to locally optimal policies that do not correspond to market equilibria. For clarity of presentation, such outcomes are treated as unsuccessful learning runs and filtered out in the numerical evaluation based on their achieved reward levels.

B. Case Study 1: Comparison with PyPSA Benchmark

This subsection compares the proposed RL-based co-optimization with a benchmark using PyPSA as a standard capacity expansion model. Continuous capacity expansion is considered for two selected lines (4-12 and 27-28). Since PyPSA does not account for strategic bidding in its transmission planning, a two-stage benchmark is considered to emulate the conventional sequential process. Specifically, we compare:

- **Two-stage benchmark** : (Stage1) Use PyPSA to compute continuous capacity expansions ΔL_ℓ under *exogenous* bidding assumptions (e.g., “high-submission” and “low-submission” regimes); (Stage2) Given those capacities, simulate the multi-agent market with RL to obtain the realized operational outcomes.
- **Proposed RL-based co-optimization**: Learn bidding policies and continuous expansion ΔL_ℓ simultaneously within the unified RL framework described in Sec. III.

We first present the results of the two-stage benchmark. The low-submission scenario assumes that all strategic generators bid at their marginal cost of \$50/MWh, while the




Fig. 2: Topology of IEEE 30-bus system

high-submission scenario assumes more aggressive bidding at \$90/MWh determined based on preliminary RL simulations. The former leads to no expansion in PyPSA, while the latter installs 81.4 MW on Line 4-12, with Line 27-28 unchanged. The resulting market outcomes obtained by running multi-agent RL under these capacities are summarized in Fig. 3 and Table I. The stacked bars in Fig. 3 present a breakdown of the total annual cost, where the colored segments represent the realized generator *revenues* (rather than *profits*) and the purple segment corresponds to the transmission *expansion cost*. Under the low-submission assumption, the lack of expansion results in severe congestion, resulting in a substantially higher realized total cost compared with the PyPSA estimate. Under the high-submission assumption, the expansion on Line 4-12 greatly reduces congestion, and the realized cost remains close to the PyPSA estimate. A slight increase is observed because, in PyPSA, the \$90/MWh bids of Gen 2 and Gen 3 prevent them from being dispatched at all, whereas in the RL simulation they sometimes bid above the \$55/MWh marginal cost of the non-strategic generators (as shown in Table I) and are dispatched at these higher prices, increasing the realized total cost.

In contrast to the two-stage benchmark, the proposed RL-based co-optimization jointly learns both the bidding strategies and the transmission expansion. The learned design results in an expansion of 76.1 MW on Line 4-12. Compared with the two-stage benchmark, particularly the high-submission case, the co-optimized solution achieves a slight reduction in expansion cost. The two-stage approach assumes uniformly high bids for all strategic generators and therefore installs more capacity than necessary, whereas the proposed method internalizes the strategic behaviors and selects a more efficient capacity level. While such integrated planning could in principle be formulated using multi-level optimization such as [3], [4], the multi-agent RL-based approach facilitates scalability to longer simulated time horizons, as the learning and simulation process can be implemented in a distributed manner across multiple workers, following the learning-based co-optimization framework in our previous work [15]. Furthermore, the current formulation can readily accommodate more

TABLE I: Summary of market simulation outcomes in case studies

Case	Scenario	Average bid price [\$/MWh]			Profit [M\$]			Operational cost [M\$]	Expansion cost [M\$]
		Gen 1	Gen 2	Gen 3	Gen 1	Gen 2	Gen 3		
1	PyPSA (low-submission)	50.0	50.0	50.0	0	0	0	103.85	0
	Two-stage (low-submission)	93.7	58.1	69.8	26.2	0.66	1.62	138.39	0
	PyPSA (high-submission)	90.0	90.0	90.0	6.17	0	0	119.64	8.14
	Two-stage (high-submission)	93.0	58.0	63.7	5.42	0.274	1.24	119.71	8.14
	Proposed co-optimization (continuous)	91.5	60.1	62.4	5.72	0.483	1.43	120.09	7.61
2	Proposed co-optimization (discrete)	75.6	94.2	84.6	4.82	0	0.0351	118.00	10.00




Fig. 3: Total cost comparison and component breakdown

complex settings such as multi-market environments.

C. Case Study: Discrete Siting Decisions

To demonstrate the applicability of the proposed framework to more realistic planning problems, we next consider a discrete siting case. The six candidate lines introduced in Sec. IV-A are included as binary expansion decisions, and the additional capacity is fixed at 50 MW for all lines. Under this setting, the design policy learns to upgrade two candidate lines 1–2 and 4–12. Although the expansion cost increases relative to the results in case 1 (Sec. IV-B), due to the additional upgrade of Line 1-2, the operational cost decreases sufficiently to yield a total cost that remains comparable to Case 1. A key factor behind this operational cost reduction is the change in Gen 1’s bidding behavior induced by the upgrade of line 1-2. In Case 1, even with the expansion of line 4-12, Gen 1 still enjoyed significant market power, and thus maintained high bid prices around \$90/MWh. By contrast, the additional capacity on line 1-2 reduces congestion on this corridor Gen 1’s ability to exert such market power. These results demonstrate that the proposed framework can identify effective expansion sites by internalizing the impact of strategic market behavior.

V. CONCLUSIONS

This paper presented an RL-based framework for co-optimization of strategic bidding and transmission expansion planning in electricity markets. By extending a multi-agent RL algorithm with a learnable continuous/discrete design policy, the proposed method enables simultaneous learning of bidding strategies and investment decisions while capturing their mutual influence. The effectiveness of the proposed framework has been shown based on case studies on the IEEE

30-bus system. The proposed framework provides a promising foundation for future applications to larger systems, multi-market environments, and integrated generation-transmission co-planning.

REFERENCES

- [1] S. Gómez and L. Olmos, “Coordination of generation and transmission expansion planning in a liberalized electricity context—coordination schemes, risk management, and modelling strategies: A review,” *Sustainable Energy Technologies and Assessments*, vol. 64, p. 103731, 2024.
- [2] E. E. Sauma and S. S. Oren, “Proactive planning and valuation of transmission investments in restructured electricity markets,” *Journal of Regulatory Economics*, vol. 30, no. 3, pp. 358–387, 2006.
- [3] E. E. Sauma and S. S. Oren, “Economic criteria for planning transmission investment in restructured electricity markets,” *IEEE Trans. Power Syst.*, vol. 22, no. 4, pp. 1394–1405, 2007.
- [4] L. P. Garces, A. J. Conejo, R. Garcia-Bertrand, and R. Romero, “A bilevel approach to transmission expansion planning within a market environment,” *IEEE Trans. Power Syst.*, vol. 24, no. 3, pp. 1513–1522, 2009.
- [5] G. Pesántez, W. Guamán, J. Córdova, M. Torres, and P. Benalcázar, “Reinforcement learning for efficient power systems planning: A review of operational and expansion strategies,” *Energies*, vol. 17, no. 9, 2024.
- [6] Y. Ye, D. Qiu, M. Sun, D. Papadaskalopoulos, and G. Strbac, “Deep reinforcement learning for strategic bidding in electricity markets,” *IEEE Trans. Smart Grid*, vol. 11, no. 2, pp. 1343–1355, 2020.
- [7] Y. Liang, C. Guo, Z. Ding, and H. Hua, “Agent-based modeling in electricity market using deep deterministic policy gradient algorithm,” *IEEE Trans. Power Syst.*, vol. 35, no. 6, pp. 4180–4192, 2020.
- [8] Y. Du, F. Li, H. Zandi, and Y. Xue, “Approximating nash equilibrium in day-ahead electricity market bidding with multi-agent deep reinforcement learning,” *J. Mod. Power Syst. Clean Energy*, vol. 9, no. 3, pp. 534–544, 2021.
- [9] W. MingKui, C. ShaoRong, Z. Quan, Z. Xu, Z. Hong, and W. YuHong, “Multi-objective transmission network expansion planning based on reinforcement learning,” in *2020 IEEE Sustainable Power and Energy Conference (iSPEC)*, pp. 2348–2353, 2020.
- [10] Y. Wang, L. Chen, H. Zhou, X. Zhou, Z. Zheng, Q. Zeng, L. Jiang, and L. Lu, “Flexible transmission network expansion planning based on dqn algorithm,” *Energies*, vol. 14, no. 7, 2021.
- [11] J. Dong, J. Cao, Y. Lu, Y. Zhang, J. Li, C. Xu, D. Zheng, and S. Han, “Transmission expansion planning: A deep learning approach,” *Sustainable Energy, Grids and Networks*, vol. 41, p. 101585, 2025.
- [12] C. Schaff, D. Yunis, A. Chakrabarti, and M. R. Walter, “Jointly learning to construct and control agents using deep reinforcement learning,” in *2019 Int. Conf. Robotics and Automation (ICRA)*, pp. 9798–9805, 2019.
- [13] T. Chen, Z. He, and M. Ciocarlie, “Hardware as policy: Mechanical and computational co-optimization using deep reinforcement learning,” in *2020 Conference on Robot Learning*, pp. 1158–1173, 2021.
- [14] M. Cauz, A. Bolland, C. Ballif, and N. Wyrsh, “Reinforcement learning for efficient design and control co-optimisation of energy systems,” in *ICML 2024 AI for Science Workshop*, p. 68, 2024.
- [15] T. Mantani, H. Hoshino, and E. Furutani, “Optimal battery sizing for real-time renewable energy bidding based on reinforcement learning,” *IEEE Trans. Energy Markets, Policy and Regulation*, 2025. (Early Access) <http://doi.org/10.1109/TEMPr.2025.3645733>.
- [16] T. Brown, J. Hörsch, and D. Schlachtberger, “PyPSA: Python for power system analysis,” *J. Open Res. Softw.*, vol. 6, no. 1, 2018.
- [17] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 2nd ed., 2018.