

Physics-Guided Machine Learning for Carbon Emission Modeling under System Disruptions: Methods and Challenges

Yihao Zhang

*School of Computing, Ulster University, Belfast, United Kingdom
Zhang-Y45@ulster.ac.uk*

Abstract. Disruptions such as pandemics and energy shocks weaken the reliability of carbon-emission models. Physics-guided machine learning (PIML) offers a practical response. This review explains how PIML improves robustness when temporal continuity breaks and system structure shifts. It covers four methods: physics-informed neural networks; hybrid or embedded designs that use CTM outputs and differentiable operators; post-processing that enforces physical feasibility; and structured or graph-based models that encode conservation and transport while keeping interpretability. Evidence shows these methods are more robust with sparse data and under distribution shifts. Yet challenges remain: balancing loss terms, dealing with complex boundaries, keeping physical inputs up to date, and the high cost of multi-scale problems. On the application side, the review focuses on pandemic-driven demand contraction, energy-supply shocks with fuel switching, and policy-induced structural change. These scenarios guide evaluation and benchmarks and help improve reliability under disruptions. On the practice side, recommended tactics include reliability-weighted fusion, adaptive coupling, explicit uncertainty handling, and dynamic graphs.

Keywords: Physics-Guided Machine Learning, Carbon Emission Modeling, Distribution Shift

1. Introduction

As countries adopt increasingly ambitious net-zero targets and face growing environmental pressures, accurate modeling and prediction of carbon emissions have become essential for energy system transition, policy design, and climate governance [1]. Carbon emission models are now widely used at national, regional, and industrial levels to assess environmental impacts, plan decarbonization pathways, and track progress toward climate goals. Under steady conditions, conventional methods ranging from process-based simulations to purely data-driven models perform well. In contrast, their performance often drops during shocks such as pandemics, energy shortages, or abrupt policy changes.

In recent years, these disruptions have become more frequent and more severe. The COVID-19 pandemic, global energy crises, extreme weather, and sudden policy shifts changed how carbon-

emitting systems operate. Such shocks weaken the assumptions of temporal continuity and structural stability used by many conventional projections [2]. During the pandemic, transport emissions fell sharply, but many industrial and residential sources remained relatively steady, so links that once held broke down. In Europe in 2022, geopolitical pressures and supply limits reduced natural gas use and led to a short-lived return to coal, which altered regional CO₂ patterns [3].

In this setting, machine-learning models such as long short-term memory (LSTM) networks and Extreme Gradient Boosting (XGBoost) face challenges with interpretability and stability. When test data differ from training data—as in building energy forecasting, deep models often lose accuracy [4]. Because they are trained mainly on historical records, they miss system-level context for disturbances and do not encode physical laws. Reliability and internal consistency therefore fall during shocks [5]. Without physical constraints or domain knowledge, outputs can be biased or drift, which weakens their value for emission-reduction planning.

To address these problems, researchers have turned to Physics-Guided Machine Learning (PIML), a set of approaches that add physical laws, scientific knowledge, and system rules to data-driven models [5, 6]. The term began as “Physics-Informed ML” and now covers both physics-guided and physics-informed designs. Studies in fluid dynamics, geophysics, and environmental modeling report better generalization and stronger physical consistency with PIML [5]. In carbon-emission modeling, using physics-based priors in machine-learning workflows can raise robustness, improve interpretability, and support policy use, especially when systems are unstable or disrupted.

This review examines PIML for carbon-emission modeling in disrupted settings. It summarizes key techniques and application areas and weighs their strengths and limits for unstable emission systems [6]. The goal is to show how physics-guided learning can improve robustness and policy relevance and to encourage links between machine learning, physical modeling, and climate science.

2. Background and research landscape

2.1. Carbon emission modeling approaches

Historically, carbon-emission modeling relied on two main routes: process-based models and data-driven machine-learning models [7]. In recent years, hybrid designs have been used to draw on both.

Process-based models simulate how emissions evolve by modeling the physics, chemistry, and engineering of energy conversion and fuel combustion. They are widely used in integrated assessment and energy system planning because of their interpretability and policy relevance [1, 2]. Their accuracy, however, depends heavily on parameter calibration and boundary assumptions, so they are sensitive to structural shifts in emission systems.

Data-driven models, such as statistical regression, support vector regression (SVM), ensemble tree methods (e.g., Random Forest, XGBoost), and deep learning time-series networks, learn patterns directly from historical records. These methods offer flexibility, scalability, and computational efficiency, especially when large observational datasets are available. A comparative study on daily CO₂ prediction found that, under non-stationary conditions, machine learning and deep learning generally outperform traditional statistical techniques at capturing nonlinear and temporal patterns [4].

Hybrid approaches combine the interpretability of process-based models with the adaptability of data-driven methods, often by using process-model outputs as features or embedding simplified physical constraints. They have been used successfully in time-series emission forecasting [8]. Most implementations, however, are ad hoc or empirical and lack a unified way to embed physical

knowledge. This gap has led to the exploration of Physics-Guided Machine Learning (PIML) as a more systematic solution for future carbon emission modeling.

2.2. Systemic disruptions and modeling challenges

Recently, carbon-emitting systems have been hit by repeated shocks: energy crises, abrupt policy changes, extreme weather, and the COVID-19 pandemic. These events changed emission structures and undercut the assumptions of temporal continuity and structural stability used by many conventional models [2]. During the pandemic, near-real-time CO₂ data made those shifts explicit [9].

Process-based models rely on fixed boundaries and tuned parameters. They struggle when transitions are abrupt. Data-driven models, in turn, degrade under non-stationarity and distribution shifts, which lowers robustness and decision value during disruptions [4]. Together, these limits point to frameworks that integrate domain knowledge with physical constraints, the basis for PIML.

2.3. Key modeling challenges and research trends

Overall, current carbon-emission modeling still struggles when systems are dynamic or perturbed. Process-based models depend on parameter tuning and static boundaries and are vulnerable to structural change. Data-driven models, meanwhile, face non-stationarity and distribution shifts and often lack clear physical interpretation [7]. Hybrid methods ease some issues but are mostly ad hoc or empirical and still lack a consistent way to embed physical knowledge [5].

Recent work supports combining physical principles and domain expertise with machine learning to improve generalization, stability, and physical consistency in complex or non-stationary settings [6]. These trends favor adopting Physics-Guided Machine Learning (PIML), discussed next.

3. Physics-Guided Machine Learning (PIML) for carbon emission modeling

3.1. PIML methodologies for carbon emission modeling

To address these limits (Section 2), Physics-Guided Machine Learning (PIML) adds physical knowledge and system constraints to the machine-learning pipeline [6]. In practice, it helps models hold up under disruption and still generalize, while preserving physical consistency and interpretability. This makes it well suited to carbon-emission modeling.

In carbon-emission modeling, PIML is usually grouped into four method categories:

(1) Physics-Informed Neural Networks (PINNs): encode governing equations or conservation laws as training signals [5]. They are used for CO₂ migration and carbon capture, where the constraints improve reliability as conditions change.

(2) Hybrid and embedded approaches: bring process-model outputs or physics-based features into the learner as structured inputs [8]. They support short-term emission forecasts and dynamic energy-system analysis and are typically more robust and physically plausible than purely data-driven baselines.

(3) Post-processing and physics-based correction: adjust predictions after training so they satisfy conservation or boundary rules. For example, CO₂ concentration correction models trained with GEOS-Chem simulations and satellite data are used for bias control [10].

(4) Structured or graph-based Methods: use graph neural networks or other structure-aware models to represent emission networks and energy flows. Examples include physics-encoded GNNs for air pollution and knowledge-guided models for soil GHG flux estimation [11, 12].

Together, these methods link mechanistic modeling with data-driven prediction and help models remain robust and physically plausible in dynamic emission systems. The following sections review each category and its representative applications.

3.2. Physics-Informed Neural Networks (PINNs)

Physics-Informed Neural Networks (PINNs), first proposed by Raissi et al., are a class of machine learning methods that add domain-specific physical knowledge, often in the form of partial differential equations, into the loss function. By applying such constraints as soft penalties, PINNs keep physical consistency during learning [5]. Unlike typical supervised learning methods that depend heavily on large labeled datasets, PINNs use governing equations to guide training. This allows them to stay robust and stable even with sparse data or under disturbed conditions.

For carbon emission modeling, PINNs work especially well when observational data are limited but the physical mechanisms are well defined. For example, recent work has used PINNs to simulate underground CO₂ storage site dynamics. By adding multiphase flow equations into the network, these models tracked pressure and saturation changes during CO₂ injection [13]. Other studies have used PINNs for inverse modeling of pollutant transport under complex wind fields, solving backward convection-diffusion problems and improving both accuracy and stability in source localization [14].

Advances have extended PINNs' reach in environmental and carbon applications. One line integrates multifidelity modeling into PINNs, combining coarse atmospheric simulations with high-resolution observations to balance accuracy and computational cost, especially for regional carbon flux estimation [15]. Another introduces adaptive weighting based on gradient normalization to balance multiple loss terms, stabilizing training under varied constraints [16]. These improvements have increased PINNs' scalability and numerical stability, enabling larger and higher-dimensional tasks.

Still, PINNs face challenges. In urban CO₂ dispersion modeling with complex boundaries, they can converge slowly and be sensitive to initialization and loss weighting [13, 14]. For multiscale carbon flux tasks, training cost grows quickly as model detail and physical complexity rise [15]. To address these issues, some studies use PINNs as modules in hybrid models, combining their strengths with other machine learning approaches. In practice, PINNs work best when data are sparse but equations are known. Scalability can be improved via multifidelity surrogates and adaptive loss weighting, and for complex boundaries or high runtime cost, by embedding PINNs in hybrid workflows [15, 16].

3.3. Hybrid and embedded approaches

Hybrid and embedded approaches are a key branch of physics-guided machine learning (PIML). Instead of enforcing governing equations in the loss function as in PINNs, these methods link process model components with statistical learners at the input, architecture, or training stages. The aim is to combine the mechanistic fidelity and interpretability of process models with the pattern-recognition ability and adaptability of machine learning, especially under non-stationary or data-scarce conditions.

In carbon emission modeling, a common approach is to feed process-model outputs, such as emission inventories, energy-balance estimates, or Chemical Transport Model (CTM) simulations, into learning frameworks as auxiliary predictors. For CO₂ flux estimation, embedding GEOS-Chem outputs in deep pipelines improves spatial generalization and uses CTM conservation and transport

relations to strengthen the physical consistency of predictions [8, 10]. Another approach extracts physics-based features from process simulations (e.g., wind vectors, stability indices, boundary-layer height, combustion emission factors) and adds them into the feature space, reducing dependence on large labeled datasets and helping avoid overfitting under distribution shifts [8].

At the architecture level, recent work has embedded lightweight mechanistic operators, such as differentiable atmospheric transport or decay modules, alongside statistical learners to preserve key processes, stabilize training, and improve physical consistency without large computational cost [17]. In regional carbon-transport tasks, these modules can raise physical plausibility without the full cost of high-resolution CTMs; combined with CTM products like GEOS-Chem inside neural networks, they can also support near-real-time CO₂ flux estimation [10]. Overall, these designs offer a balance between physical fidelity and computational efficiency for operational emission modeling [17, 8].

Hybrid and embedded designs tend to be more robust under system disruptions: physics-based priors provide stable constraints that reduce drift when distributions shift, while the data-driven parts adapt to changed regimes [8]. However, performance depends on the accuracy, timeliness, and coverage of the physical components; bias or delay in process model outputs can spread through the pipeline [8]. To keep resilience under disruptions, use reliability-weighted feature fusion, adaptive coupling, and near-real-time model and data integration [8]. In short, hybrid and embedded designs work best with reliable process-model products or physics-based features and operational timelines; robustness under shift benefits from reliability-weighted fusion and adaptive coupling [8].

3.4. Post-processing for physical consistency

Unlike training-stage coupling of physical priors (e.g., PINNs in Section 4.2 and hybrid schemes in Section 4.3), post-processing and physics-based corrections are applied after predictions are made. They improve results by enforcing feasibility and physical coherence without retraining the upstream model. Common methods include:

- (1) Constraint projection or feasible-set mapping, which maps outputs back to sets that meet non-negativity, sectoral and regional budget consistency, and mass-balance closure [18];
- (2) CTM-guided residual correction, which uses CTM (e.g., GEOS-Chem) backgrounds, transport footprints, and diagnostics as physical anchors to learn residuals from biases and correct them [10];
- (3) Multi-source reconciliation, aligning satellite retrievals and in-situ observations with priors to reduce systematic drift [8].

For disrupted systems, post-processing can be a low-latency final step: physical constraints stabilize outputs when statistical patterns drift, and external model or observation cues help suppress bias [8]. Effectiveness depends on the timeliness and reliability of both the upstream predictor and the physical drivers; delays or biases in CTM inputs can propagate errors. In practice, use uncertainty-aware weighting, adaptive gains, and near-real-time integration with open diagnostics (e.g., mass-balance residuals, budget-closure errors) to keep robustness [18].

These methods are easy to integrate, but their performance depends on reliable and timely upstream data. Delays or biases can cause error cascades. Rigid feasibility projections may hide real regime changes, and residual corrections often address surface bias rather than deeper model issues. Uncertainty handling is often implicit. We suggest pairing uncertainty-aware gains and thresholds with sensitivity checks to prevent drift [8], and applying conservation- and budget-based diagnostics for monitoring and correction [18]. Operationally, treat post-processing as a quick final step to

enforce feasibility under disruptions, with continuous monitoring via conservation and budget diagnostics [18].

3.5. Structured and graph-based PIML

In structured and graph-based methods, emission systems are modeled as networks: nodes are sources, sinks, sectors, or grid cells, and edges capture transport, energy flows, or supply-chain links. This setup allows physics-aware architectures to encode conservation laws, node-level mass and energy balance, and directed transport or diffusion as inductive biases, while keeping interpretability at the subsystem level.

A common case is the physics-encoded graph neural network. Designs usually combine: (i) guiding message passing with domain structure and external fields; (ii) selecting physically relevant node and edge features (e.g., emission intensity, fuel mix, wind vectors, stability indices, boundary layer height); and (iii) building adjacency and edge weights from transport sensitivities or engineered connectivity. Together, these improve generalization and physical plausibility in atmospheric modeling [11]. Extending this, knowledge-guided graph models add process-model knowledge, simulation priors, and multi-graph structures, enabling flux estimation in data-scarce regimes [12].

These models can work alone or inside larger PIML pipelines, for example, using CTM-based connectivity for input coupling, or applying graph-regularized projection in post-processing to ensure feasible aggregation at sector or region level. Limits include sensitivity to graph design, mismatches between graph and observation scales, and the cost of long-horizon spatiotemporal inference. In practice, when processes have clear network structure and node/edge interpretability is important, graph methods are a natural fit. Match graph resolution to observations and regularize with node balance and advection-diffusion operators. Under disruptions, dynamic graphs can update connectivity and weights with changing weather or sector activity; with node-balance and advection-diffusion regularization, they help keep physical consistency and robustness under shifts [11, 12].

4. Applications for carbon emission modeling under system disruptions

Unified evaluation plan (applies to Sections 4.1–4.3) To ensure comparability across disruption scenarios, we standardize the evaluation before discussing method selection. Unless noted otherwise, Sections 5.1–5.3 follow a change-point aware split into pre-disruption, disruption, and recovery; assess temporal generalization on held-out phases and structural generalization on held-out regions or network configurations; track physics consistency via mass-balance residuals and budget-closure errors; check uncertainty calibration using empirical interval coverage; and record update latency (from data arrival to release) to assess near-real-time usability [10, 18]. Scenario subsections add only diagnostics specific to the context.

4.1. Demand contraction under pandemic disruptions

Scenario overview: Mobility and electricity demand drop sharply; sectoral correlations decouple; operating regimes shift within short windows. This pattern was evident during COVID-19: activity contracted and cross-sector relations temporarily broke down. Models trained on pre-disruption data tend to drift and can distort sectoral or regional aggregates [9, 4].

Method selection (applicability and role in this scenario; method foundations in Section 4):

Hybrid and embedded: When recent observations are scarce and fast updates are needed, add inventories, energy balance estimates, CTM backgrounds, and mobility indicators as structured priors to stabilize short-term updates and provide physical cues [8, 10].

Post-processing: After producing initial predictions, perform bias correction and consistency checks; where appropriate, add simple conservation checks to keep sectoral and regional aggregates stable in near-real-time use [18, 10].

Structured and graph-based: When node- and edge-level attribution and network representation are required, for example, model the activity–energy–emissions system as a network to support structured inference [2, 11].

PINNs (conditional): For local subproblems where transport physics is clear but observations are sparse, use PINNs as modular components within hybrid pipelines [13].

Evaluation plan: Follow the unified evaluation plan. In this scenario, also examine how mobility-related input delays and coverage gaps affect temporal generalization, structural generalization, and update latency.

4.2. Energy supply shocks and fuel switching

Scenario overview: When imports are constrained, fuel prices surge, or generation units are suddenly offline, the power and industrial sectors adjust the fuel mix within short windows (e.g., less gas and temporary coal rebound), which shifts emission intensities and their spatiotemporal patterns. The 2022 European energy crisis is a representative case: gas use declined while coal temporarily increased, reshaping regional CO₂ trajectories [3]. For clarity, the following discussion is illustrative of mechanisms and channels; it is not intended to generalize to all regions or to provide country-level quantification.

Method selection (applicability and role in this scenario; method foundations in Section 4):

Hybrid and embedded: For rapid rolling updates with limited recent observations, use inventories, energy-balance estimates, CTM backgrounds, and indicators of supply constraints and price movements as priors or features to stabilize responses to fuel-switching-induced distribution shifts [8].

Post-processing: After the initial forecast, apply bias correction and feasibility checks; when needed, enforce simplified constraints such as power balance and carbon budget so sector- and region-level aggregates remain feasible and stable under large fuel-mix changes [18].

Structured and graph-based: Model generation assets and regions as nodes, and transmission and cross-regional interactions as edges. Use network connectivity and direction constraints to describe substitution pathways and cross-regional flows, improving attribution and extrapolation [11].

PINNs (conditional): For local subproblems with clear transport physics but sparse observations (e.g., local transport processes or other physically constrained steps), embed PINNs as modular components within hybrid pipelines to reinforce physical consistency [13].

Evaluation plan: Follow the unified evaluation plan stated at the start of this section. In this scenario, additionally track deviations in power-balance and carbon-budget closure during shock and adaptation phases, and assess how cross-regional transfer constraints from fuel switching affect structural generalization. When necessary, run near-real-time consistency checks to ensure operability.

4.3. Policy-induced structural change

Scenario overview: When multiple policy instruments take effect, such as allowance adjustments, carbon-price changes, subsidy or tax reforms, mandatory retirements, or production caps, activities and technologies are reallocated across sectors and regions within short windows, which alters emission intensities and spatial patterns [1, 2]. For readability, the discussion below explains mechanisms and usage rather than country-level magnitudes; as an illustrative case, see the 2022 European energy shock [3].

Method selection (applicability and role in this scenario; method foundations in Section 4):

Hybrid and embedded: After a policy switch, use inventories, energy-balance estimates, and CTM backgrounds as stabilizers, and include policy variables (e.g., allowances, carbon prices, sector restrictions) and activity indicators as priors or features to mitigate distribution shift and improve extrapolation robustness [8].

Post-processing: After the initial forecast, apply bias correction and feasibility checks; when needed, enforce non-negativity, sectoral and regional budget consistency, allowance limits, and simplified conservation checks. During structural breaks, align residuals using CTM backgrounds and diagnostics so that aggregates remain stable and consistent [18, 10].

Structured and graph-based: Represent sectors, regions, and key assets as nodes, and inter-regional transmission and energy exchanges as edges. Use connectivity and directionality constraints to capture policy-driven relocation paths and cross-regional flows, improving attribution and out-of-structure generalization [11].

PINNs (conditional): Use as modular components for local subproblems where physics is clear but observations are sparse (e.g., specific physically constrained steps or local transport). To ease scaling and complex boundaries, combine multifidelity strategies and adaptive loss weighting [15, 16].

Evaluation plan: Follow the unified evaluation plan stated at the start of Section 5. For this scenario, additionally track the time paths of deviations in allowance consistency and sectoral or regional carbon-budget closure during policy switch and adaptation phases; test out-of-structure generalization across regions subject to different policies; and set a simple counterfactual with and without the policy to check robustness and update latency.

5. Recommendations and outlook

Building on Section 4 (methods) and Section 5 (scenarios), this section offers operational guidance for deploying PIML in disrupted carbon emission systems. First, design disruption-aware workflows: once regime shifts are detected, pair fast learners with conservation and budget checks so near-real-time updates remain physically feasible; track mass-balance residuals and budget-closure errors as routine operational diagnostics [18]. Second, prioritize multi-source, near-real-time assimilation centered on a CTM: use CTM backgrounds and sensitivities as anchors and fuse satellite and in-situ observations to stabilize short-window estimates [10]; reuse the rapid CO₂ monitoring infrastructure and practices developed during COVID-19 [9]. Third, invest in structure-aware representations and modular physics: represent the activity-energy-emissions system as a graph to enable node- and edge-level attribution and cross-regional extrapolation [11]; deploy PINN components only where local physics is clear but observations are sparse, and pair them with multifidelity training and adaptive loss weighting to improve scalability and robustness [15, 16]. Fourth, establish disruption-aware benchmarks: release datasets aligned with process priors, share change-point partitions, and standardize diagnostics to enable comparable assessments across

methods and regions [6, 18]. Fifth, strengthen policy relevance: stress-test models under allowance changes, fuel-switching constraints, and demand collapses, and report interpretable indicators aligned with mitigation pathways to support decisions [1, 2]. Collectively, these priorities aim to maintain robustness under non-stationarity without sacrificing physical consistency or interpretability.

6. Conclusion

This article reviews physics-guided machine learning (PIML) methods for carbon-emission modeling under system disruptions. Prior surveys show that with distribution shift or sparse observations, purely data-driven models often lose generalization and physical consistency, so adding physics and system knowledge is needed. Among the four families: PINNs fit tasks with clear physics but limited data, putting conservation and governing equations into training to boost interpretability and robustness; hybrid/embedded designs pair process-model outputs and physics-based features with learners to support robust rolling updates when recent data are scarce ; post-processing methods avoid retraining by enforcing feasibility and conservation to keep aggregate totals stable during regime shifts; structured/graph methods model the activity-energy-emissions system as a network, enabling network-aware constraints and interpretable node- and edge-level inference. Looking ahead, use disruption-aware workflows: once a regime shift is detected, pair fast learners with routine checks such as conservation and budget-closure tests, and fuse multi-source data in near real time using a chemistry-transport model (CTM) as the reference to stabilize short-window estimates. Also use consistent change-point splits and consistency checks so results are comparable across regions and methods and useful for policy. These steps help keep robustness under nonstationary conditions without giving up physical consistency or interpretability.

References

- [1] International Energy Agency. (2021). Net zero by 2050: A roadmap for the global energy sector. Paris, France: IEA. Retrieved from <https://www.iea.org/reports/net-zero-by-2050>
- [2] Intergovernmental Panel on Climate Change. (2022). Climate change 2022: Mitigation of climate change—Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK: Cambridge University Press.
- [3] International Energy Agency. (2023). CO₂ emissions in 2022 – Analysis. Paris, France: IEA. Retrieved from <https://www.iea.org/reports/co2-emissions-in-2022>
- [4] Liang, X., Liu, Z., Wang, J., Jin, X., & Du, Z. (2023). Uncertainty quantification-based robust deep learning for building energy systems considering distribution shift problem. *Applied Energy*, 337, 120889.
- [5] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707.
- [6] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3, 422–440.
- [7] Zhao, C., Zhang, M., Bai, J., Wu, J., & Chang, I.-S. (2025). A review of the application of machine learning in carbon emission assessment studies: Prediction optimization and driving factor selection. *Science of the Total Environment*, 987, 179678.
- [8] Jin, Y., Sharifi, A., Li, Z., Chen, S., Zeng, S., & Zhao, S. (2024). Carbon emission prediction models: A review. *Science of the Total Environment*, 927, 172319.
- [9] Liu, Z., et al. (2020). Near-real-time monitoring of global CO₂ emissions reveals the effects of the COVID-19 pandemic. *Nature Communications*, 11, 5172.
- [10] Wang, H., Li, D., Zhou, R., Hu, X., Wang, L., & Zhang, L. (2024). A new method for top-down inversion estimation of carbon dioxide flux based on deep learning. *Remote Sensing*, 16(19), 3694.

- [11] Li, L., et al. (2023). Improving air quality assessment using physics-inspired deep graph learning. *npj Climate and Atmospheric Science*, 6, 152.
- [12] Zhang, Y., et al. (2025). KG-FGNN: Knowledge-guided GNN foundation model for fertilisation-oriented soil GHG flux prediction. *arXiv preprint arXiv: 2506.15896*.
- [13] Shokouhi, P., Kumar, V., Prathipati, S., Hosseini, S. A., Giles, C. L., & Kifer, D. (2021). Physics-informed deep learning for prediction of CO₂ storage site response. *Journal of Contaminant Hydrology*, 241, 103835.
- [14] Chuprov, I., Derkach, D., Kychkin, A., & Efremenko, D. (2025). Application of physics-informed neural networks for solving the inverse advection–diffusion problem to localize pollution sources. *arXiv preprint arXiv: 2503.18849*.
- [15] Penwarden, M., Zhe, S., Narayan, A., & Kirby, R. M. (2022). Multifidelity modeling for physics-informed neural networks (PINNs). *Journal of Computational Physics*, 451, 110844.
- [16] Zhu, P., Liu, Z., Xu, Z., & Lv, J. (2025). An adaptive weight physics-informed neural network for vortex-induced vibration problems. *Buildings*, 15(9), 1533.
- [17] Beucler, T., et al. (2021). Enforcing analytic constraints in neural networks emulating physical systems. *Physical Review Letters*, 126, 098302.
- [18] Weir, B., Ott, L. E., Collatz, G. J., Kawa, S. R., Poulter, B., Chatterjee, A., Oda, T., & Pawson, S. (2021). Bias-correcting carbon fluxes derived from land-surface satellite data for retrospective and near-real-time assimilation systems. *Atmospheric Chemistry and Physics*, 21, 9609–9628.