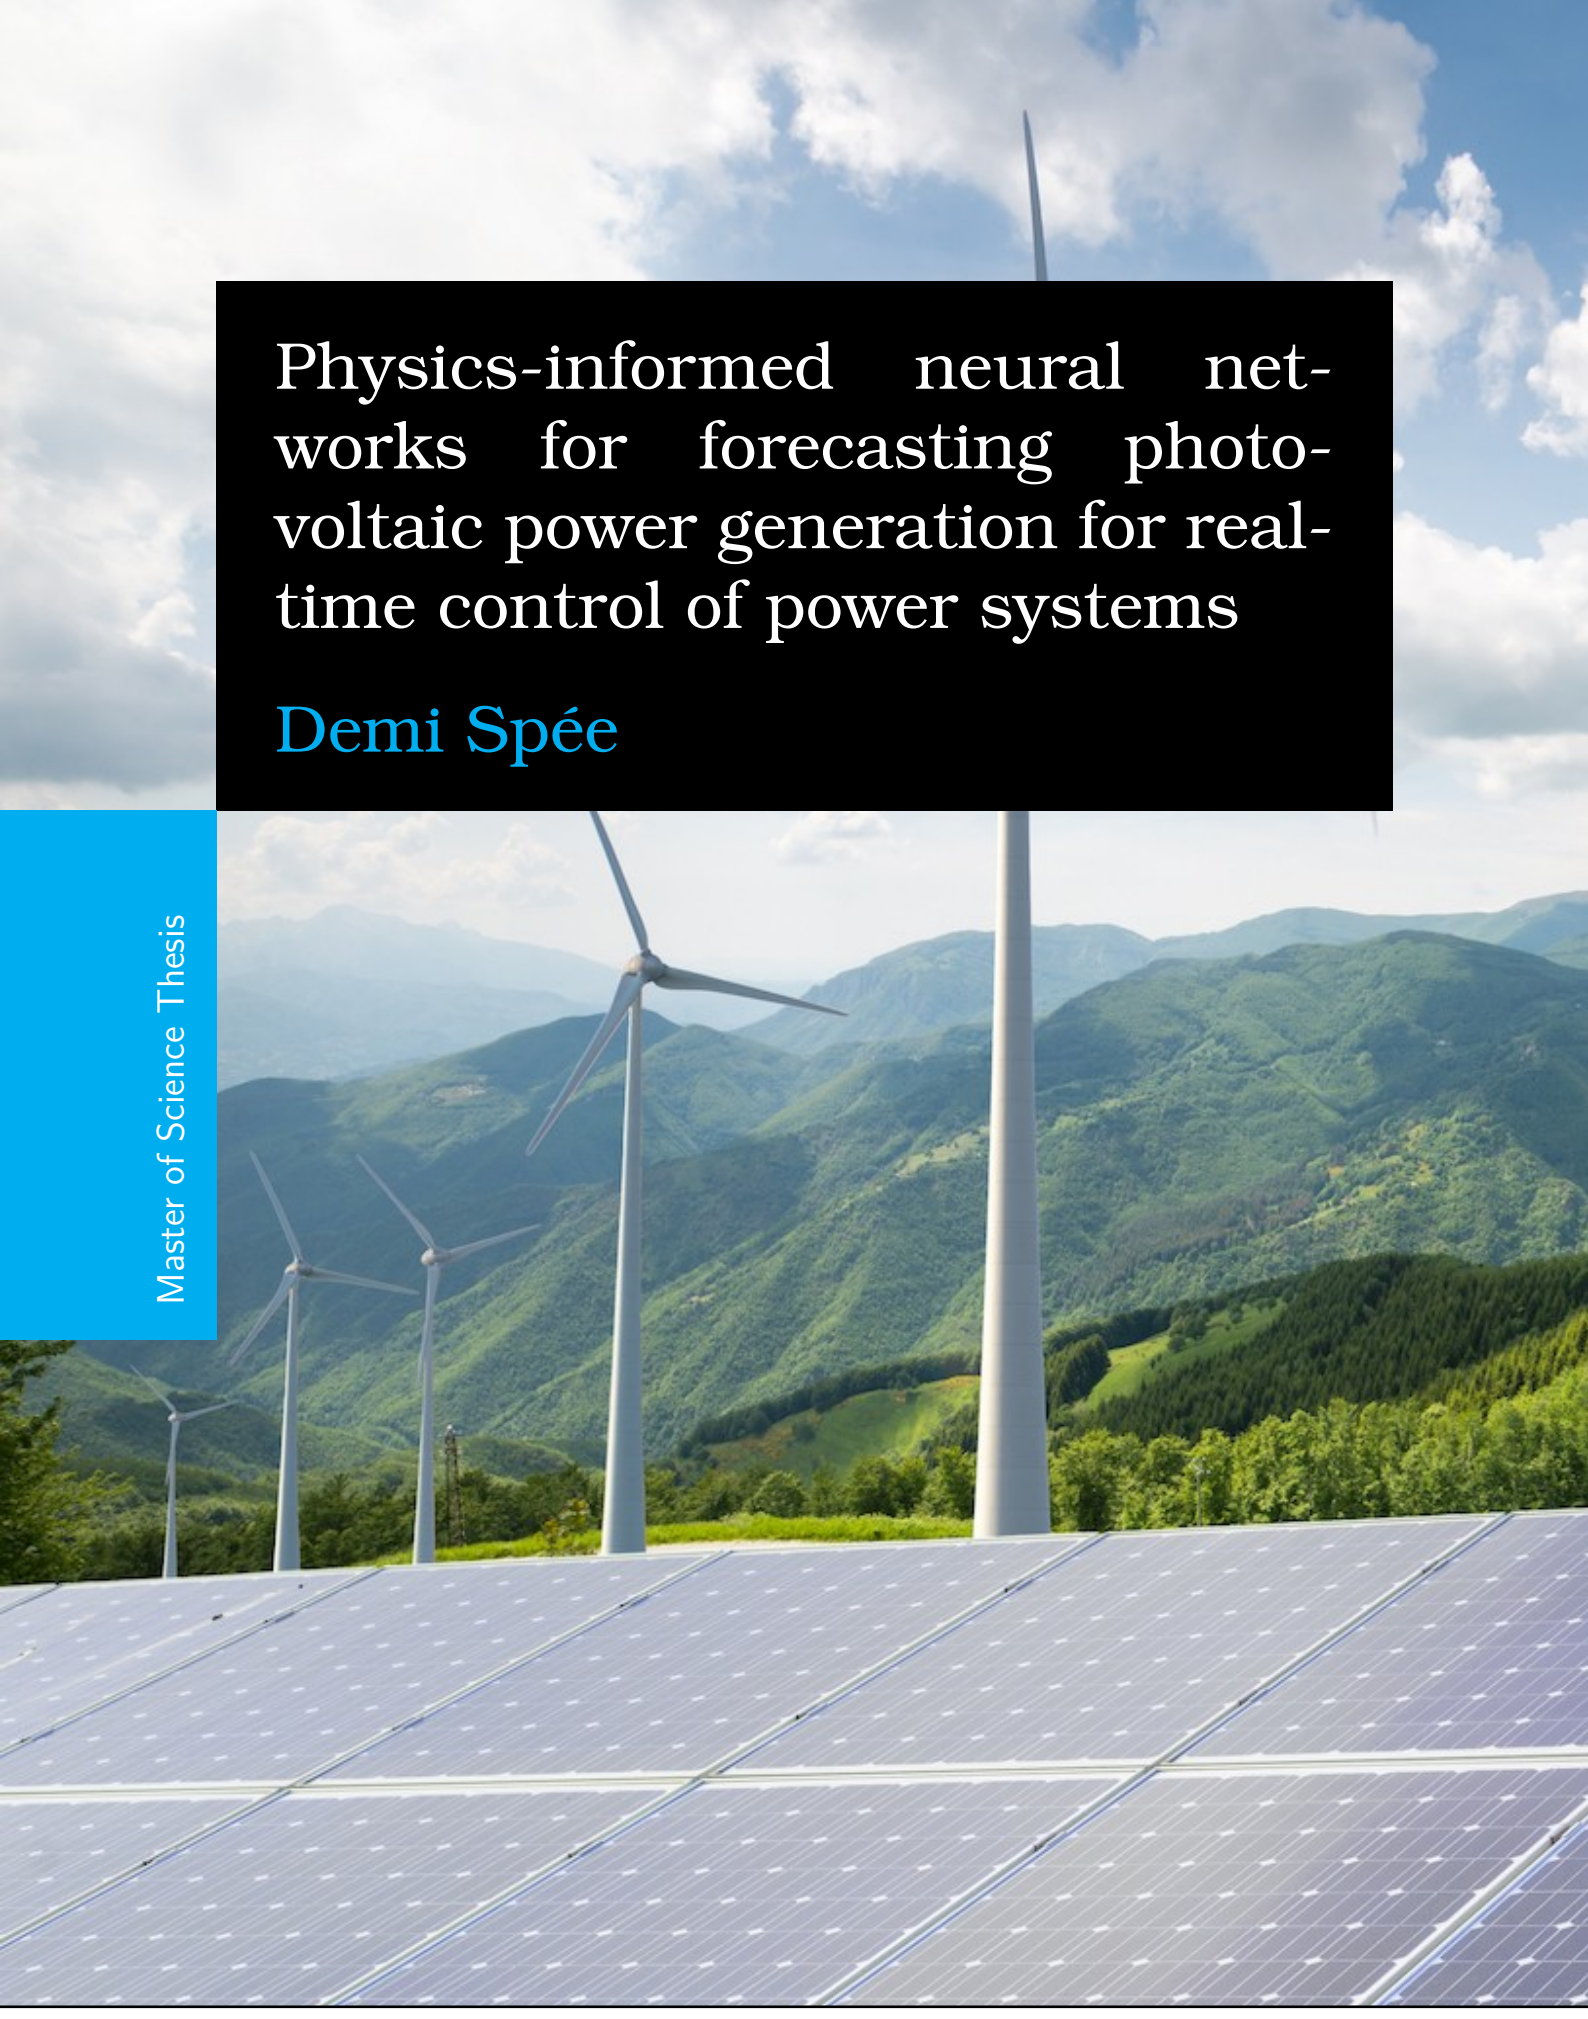

Physics-informed neural networks for forecasting photovoltaic power generation for real-time control of power systems

Demi Spée

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

Physics-informed neural networks for forecasting photovoltaic power generation for real-time control of power systems

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

Demi Spée

January 12, 2026

Faculty of Mechanical Engineering (ME) · Delft University of Technology

Abstract

Modern power systems are being fundamentally reshaped by the increasing integration of renewable energy sources (RESs) [48], such as solar and wind power generation. Their inherent variability poses significant challenges for real-time grid control. In particular, advanced control strategies, including model predictive control, require multi-step forecasts at temporal resolutions relevant to the specific control application, which in this work is of the order of seconds. However, commonly available meteorological datasets are often temporally sparse, creating a pronounced mismatch between data availability and control requirements. To address this challenge, this thesis investigates the use of physics-informed neural networks (PINNs) to forecast photovoltaic (PV) power generation in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity.

The PINN developed incorporates partial physical knowledge of atmospheric processes through a spatiotemporal cloud motion equation, enabling physically plausible generalization to unseen time steps and conditions. To generate multi-step forecasts at temporal resolutions on the order of seconds, a recursive forecasting framework is developed. In addition, an improved training strategy is proposed in which the model is trained using its own recursively generated forecasts, thereby reducing the mismatch between training and inference. Finally, a novel correlation-based collocation point sampling strategy is developed to generate physically plausible and statistically representative collocation points, thereby supporting effective physics-based regularization.

The proposed methods are evaluated through two complementary case studies. In a case study based on a selected region in France, the PINN with the improved training strategy demonstrates a clear advantage under challenging atmospheric conditions compared to a persistence model, which assumes constant behavior over the forecast horizon. It also achieves superior data efficiency relative to purely data-driven neural networks (NNs) and improved computational efficiency over a physics-only model during inference. The proposed collocation point sampling strategy also consistently outperforms uniform sampling.

In addition, a real-world case study based on a measurement site in Hawaii validates the forecasting framework under a realistic PV deployment scenario in which available data are not only temporally sparse but also satellite-derived. Because such data provide only an

approximate, temporally and spatially smoothed representation of the true atmospheric conditions affecting the PV system, this setting introduces an additional data-reality mismatch. Nevertheless, the framework remains capable of producing meaningful forecasts at temporal resolutions on the order of seconds, establishing its practical applicability for PV systems worldwide without local measurement infrastructure.

Table of Contents

Acknowledgements v

1 Introduction 1
- 1-1 General introduction . . . 1
- 1-2 Problem description . . . 2
- 1-3 Research questions . . . 3
- 1-4 Outline of the thesis . . . 5

2 Background and related literature 7
- 2-1 Photovoltaic power generation forecasting . . . 7
  - 2-1-1 Power system control requirements . . . 7
  - 2-1-2 Classification of forecasting methods . . . 8
- 2-2 Physics-informed neural networks (PINNs) . . . 9
  - 2-2-1 Formulation of the PINN framework . . . 9
  - 2-2-2 Collocation point sampling . . . 10
- 2-3 Multi-step and recursive forecasting . . . 11
- 2-4 Summary . . . 12

3 Methodology of the PINN-based forecasting framework 13
- 3-1 PINN design . . . 13
  - 3-1-1 Forecast target and model inputs . . . 14
  - 3-1-2 Physical model: cloud motion . . . 15
  - 3-1-3 Training loss function . . . 15
- 3-2 Collocation point sampling strategy . . . 17
  - 3-2-1 Correlation analysis . . . 17
  - 3-2-2 Correlation-based input variable clustering . . . 18

3-2-3 Cluster-wise joint sampling . . . 18
3-2-4 Temporal and spatial sampling . . . 19
3-3 Normalization and scaling of inputs and governing equation . . . 20
3-3-1 Input normalization . . . 20
3-3-2 Governing equation scaling . . . 21
3-4 Recursive forecasting framework . . . 21
3-4-1 Forecasting procedure during inference . . . 21
3-4-2 Forecasting procedure for validation . . . 23
3-5 Improved training strategy using recursive forecasting . . . 23
3-5-1 Training procedure . . . 23
3-5-2 Discussion . . . 24
3-6 Comparison models . . . 24
3-6-1 Persistence model . . . 25
3-6-2 Data-driven neural network models . . . 25
3-6-3 Physics-only cloud motion model . . . 25
3-7 Validation design . . . 26
3-7-1 Reproducibility . . . 26
3-7-2 Evaluation metrics . . . 26
3-8 Summary . . . 27

4 Case study **29**
4-1 Experimental setups . . . 29
4-2 Data sources and preprocessing . . . 30
4-2-1 Solcast coarse-resolution satellite-derived data . . . 30
4-2-2 NREL high-resolution measurement data . . . 32
4-3 Model architecture and experimental configuration . . . 32
4-4 Simulation results for the France case study . . . 34
4-4-1 Baseline PINN: forecast behavior and sensitivity analysis . . . 36
4-4-2 Improved PINN: forecast behavior and sensitivity analysis . . . 41
4-5 Simulation results for the Hawaii case study . . . 45
4-6 Discussion . . . 47
4-7 Summary . . . 48

5 Conclusions and recommendations **51**
5-1 Contributions . . . 51
5-2 Conclusions . . . 52
5-3 Recommendations for future work . . . 53

A Paper **55**

Glossary **71**

Acknowledgements

I would like to sincerely thank Prof. Dr. Ir. Bart De Schutter for giving me the opportunity to carry out this Master's thesis within his research group and for his guidance throughout the project.

I also would like to thank my daily supervisors, Ir. Francesco Cordiano and Ir. Alessandro Riccardi, for their continuous support and involvement throughout the entire thesis process. They guided me through all stages of the thesis, with weekly meetings to discuss progress, solve problems together, and explore new directions. They consistently provided clear, critical, and detailed feedback, often beyond regular working hours, helping me bring my work to a higher level. I sincerely thank them for their encouragement and insightful suggestions, which made working on this thesis not only a great learning experience but also a very rewarding one.

Delft, University of Technology — Demi Spée
January 12, 2026

Chapter 1

Introduction

1-1 General introduction

The global transition from centralized, dispatchable fossil-fueled power generation to decentralized, weather-dependent renewable energy sources (RESs) is fundamentally reshaping modern power systems [48, 61]. This transition is driven by the urgency to mitigate climate change and meet international decarbonization goals, and promises a more sustainable and environmentally responsible energy future [28]. At the same time, it introduces new challenges for the reliable operation of power systems, mainly due to the inherent variability, limited controllability, and decentralized nature of RESs [39, 77].

In Europe, the share of renewable energy in total energy consumption reached a historical high of 24.5% in 2023 [28], as shown in Figure 1-1, driven mainly by a substantial growth of solar and wind power generation. However, reaching the newly adopted European Union target of at least 42.5% by 2030 will require more than doubling the current rate of RESs deployment, further increasing the influence of weather-dependent power generation on power system operation.

Reliable electricity supply is a cornerstone of modern society, supporting a wide range of critical services and economic activities in sectors such as healthcare, transportation, communication, and industry [57]. Power systems must therefore balance electricity generation and consumption to maintain a stable grid frequency close to its nominal value. Imbalances can cause frequency deviations that disrupt industrial processes, damage equipment, and in severe cases trigger cascading failures that lead to large-scale blackouts [36]. Recent events, such as the April 2025 blackout in the Iberian Peninsula, have highlighted the challenges of maintaining frequency stability in power systems with a high share of RESs [16].

Historically, frequency stability has been maintained through the real-time control of large, centralized, dispatchable generators [48], typically referred to as load frequency control (LFC). As these generators are increasingly displaced or complemented by non-dispatchable RES, traditional control paradigms face growing challenges [4]. In this context, accurate and reliable forecasting of renewable power generation has become increasingly important in modern power

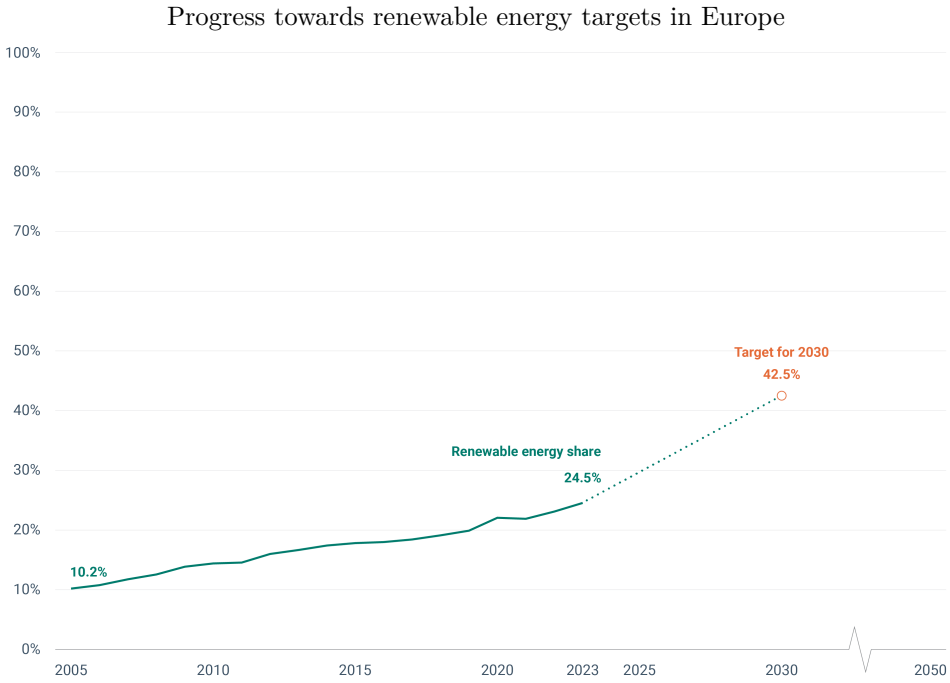

Figure 1-1: Share of renewable energy in Europe's total energy mix from 2005 to 2023 [28].

system operation, as it allows anticipating power imbalances and supports proactive control [25, 55].

A wide range of forecasting approaches have been developed to support power system operation. Broadly, these approaches can be categorized into physics-based models, which rely on physical laws describing the process of interest, and data-driven machine learning (ML) models, which can learn complex nonlinear relationships directly from historical data [9, 79]. More recently, physics-informed machine learning (PIML) approaches have emerged as a promising research direction. They combine physical knowledge with data-driven techniques, aiming to leverage the physical plausibility of physics-based models and the universal approximation capabilities of ML methods [42]. The effectiveness of both forecasting approaches in power system operation depends on several factors, including data availability, the complexity of the underlying physical processes, and the temporal and spatial resolution required by control applications.

1-2 Problem description

Although renewable power generation forecasting is widely recognized as a key enabler of reliable operation of modern power systems [55, 61], significant challenges remain in aligning the capabilities of existing forecasting methods with the requirements of real-time control applications [22]. Advanced control strategies such as model predictive control (MPC) [27, 72], which are increasingly considered for tasks such as LFC [5], require multi-step forecasts at temporal resolutions relevant to the specific control application. In this thesis, the term *control-relevant* is used to denote multi-step forecasts at temporal resolutions on the order of

seconds, corresponding to the time scales of the control layers under consideration. Meeting these requirements remains challenging.

Physics-based forecasting models can provide physically plausible forecasts, but typically require detailed knowledge of the underlying physical laws and system-specific parameters, and are computationally expensive due to the repeated numerical evaluation of the physical model [50]. However, due to the complexity of the atmospheric processes governing renewable power generation, complete physical descriptions are often unavailable or impractical [10].

On the other hand, data-driven ML models can learn complex nonlinear relationships directly from data, but require large volumes of training data at a temporal resolution that is sufficiently high for the intended forecasting task [14, 50]. However, in practice, commonly available meteorological datasets, including globally available satellite-derived services such as Solcast [80], typically provide data at 5-minute resolution. This creates a pronounced mismatch between temporally sparse observational data and the higher temporal resolution required by power system control applications. Such a mismatch is especially relevant for renewable power generation systems that lack local sensors and must rely exclusively on temporally sparse satellite-based data.

The limitations of physics-based and data-driven forecasting approaches at temporal resolutions on the order of seconds motivate the exploration of PIML frameworks [42]. Such frameworks can potentially exploit partial physical knowledge to support data-driven learning under temporal data sparsity. Within this framework, physics-informed neural networks (PINNs) have received significant attention, as they allow physical knowledge to be incorporated into the training of neural networks (NNs) through the enforcement of governing equations [68]. In principle, this enables PINNs to generate physically plausible forecasts at higher temporal resolution than that available in the data, by generalizing between temporally sparse observations. However, the practical effectiveness of PINNs in generating multi-step renewable power generation forecasts at temporal resolutions on the order of seconds under temporal data sparsity and partial physical knowledge remains an open research question.

1-3 Research questions

The problem description in Section 1-2 highlights a fundamental gap between the requirements of real-time power system control applications and the capabilities of existing renewable power generation forecasting approaches. On the one hand, advanced control strategies such as MPC require multi-step forecasts at temporal resolutions on the order of seconds. On the other hand, both physics-based and data-driven forecasting methods face intrinsic limitations in such settings, arising from partial physical knowledge of relevant atmospheric processes and temporal data sparsity, respectively. This motivates the exploration of PIML approaches as a potential means to bridge this gap.

Although these challenges apply broadly to weather-dependent RESs, this thesis focuses specifically on photovoltaic (PV) power generation. PV systems are sensitive to atmospheric conditions, where variations in power generation are primarily driven by spatiotemporal cloud motion that can induce rapid fluctuations in solar irradiance over time and space [22]. In addition, PV installations are often deployed in spatially distributed configurations with limited

local sensing, which motivates the use of temporally sparse satellite-derived data for forecasting. These characteristics make PV power generation in spatially distributed systems a representative case for investigating control-relevant forecasting under partial physical knowledge of relevant atmospheric processes and temporal data sparsity.

Against this background, the central research question of this thesis is formulated as follows:

> **How can physics-informed neural networks be used to forecast PV power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements?**

To address this overarching question in a structured and systematic way, several more specific research questions are considered. First, the effectiveness of PINNs depends on the suitability and formulation of the physical knowledge incorporated into the training process. This motivates the following research question:

- **RQ1:** What partial physical knowledge of atmospheric processes, expressed through governing equations, is suitable for PV power generation forecasting in spatially distributed PV systems, and how can it be incorporated into a physics-informed neural network?

Second, real-time power system control requires forecasts not only at relatively high temporal resolution but also over multiple future time steps. Such multi-step forecasts are typically generated recursively and are therefore susceptible to error accumulation, a challenge that becomes even more pronounced under temporal data sparsity. This motivates the following research question:

- **RQ2:** How can physics-informed neural networks be used within a recursive forecasting framework to generate physically plausible multi-step PV power generation forecasts at temporal resolutions on the order of seconds under temporal data sparsity?

Finally, the practical performance of PINNs depends on how they are designed and trained, particularly under recursive use and temporal data sparsity. This motivates an investigation into methodological adaptations that improve the forecast performance in such settings:

- **RQ3:** How can the design and training of physics-informed neural networks be adapted to improve PV power generation forecasting performance under recursive use and temporal data sparsity?

Together, these research questions define the scope of this thesis. Rather than directly addressing power system control itself, the focus is on the forecasting problem that underlies control performance. This enables a detailed and critical investigation of PINNs for PV power generation forecasting under the combined constraints of control-relevant forecasting requirements, temporal data sparsity, and partial physical knowledge. The resulting insights provide a basis for assessing the suitability of PINNs as forecasting models for real-time power system control applications operating under such practical conditions.

1-4 Outline of the thesis

This thesis is structured as follows. Chapter 2 reviews the relevant literature on PV power generation forecasting, PINNs, cloud motion modeling, and recursive multi-step forecasting, thereby positioning the present work within the existing body of research. Chapter 3 presents the methodological framework of this thesis. It describes the design of the PINN for PV power generation forecasting, the developed correlation-based collocation point sampling strategy, the developed recursive forecasting framework, and the proposed improved training strategy in which the model is trained using its own recursively generated forecasts. The chapter also introduces the comparison models and the validation design. Chapter 4 evaluates the proposed methods through two complementary case studies. The France case study provides an experimental setting for systematic evaluation and comparison of forecasting performance, including sensitivity analyses. The Hawaii case study offers an end-to-end validation for real PV systems under a realistic deployment scenario. The chapter concludes with a comprehensive discussion and summary of the experimental results. Chapter 5 summarizes the main contributions of the thesis, presents the conclusions with respect to the research questions, and outlines directions for future research.

Chapter 2

Background and related literature

This chapter reviews the background and relevant literature on photovoltaic (PV) power generation forecasting for real-time power system control. It establishes the theoretical and methodological context of the present work and highlights the key challenges that motivate the proposed approach.

The remainder of this chapter is structured as follows. Section 2-1 reviews PV power generation forecasting, including power system control requirements, common forecast targets, and the main classes of existing forecasting methods. Section 2-2 introduces physics-informed neural networks (PINNs), covering the formal formulation of the PINN framework and current approaches for collocation point sampling. Finally, Section 2-3 discusses recursive forecasting and highlights the fundamental challenges that motivate the research addressed in this thesis.

2-1 Photovoltaic power generation forecasting

This section outlines the main considerations in forecasting PV power generation, including the requirements of power system control, key design choices in defining the forecast target, and the main modeling approaches used in the literature.

2-1-1 Power system control requirements

The weather-dependent nature of PV systems introduces significant variability and uncertainty in modern power systems [77]. These disturbances occur across a wide range of time scales and must be controlled to maintain stable operation. Accordingly, power system control is organized into multiple layers that operate at different time scales. Fast control layers act on the order of seconds (e.g., LFC), while slower layers act on minutes (e.g., tertiary control) and hours or more (e.g., scheduling and dispatch) [61]. This thesis focuses on control layers that act on the order of seconds, where PV power generation is strongly influenced by rapid and complex atmospheric dynamics, particularly spatiotemporal cloud motion and the resulting modulation of solar irradiance [22].

In these fast control layers, traditional control strategies typically react to real-time deviations in system states such as frequency and voltage. In contrast, more advanced approaches, particularly predictive strategies such as MPC, explicitly incorporate multi-step forecasts of future deviations in PV power generation [27], enabling proactive control. As a result, accurate PV power generation forecasting has become an important component of modern power system operation [2].

The effectiveness of predictive control depends on the quality of the forecasts [64]. Forecast errors can translate directly into degraded control performance, including delayed frequency restoration, increased oscillations, and higher actuator usage. In addition, the physical plausibility of multi-step forecasts is also important, as physically implausible forecasts can mislead the controller and potentially compromise system stability or operational safety [25].

Forecasting PV power generation can be formulated in different ways, depending on the intended application. The first important distinction concerns the temporal resolution of the forecasts. Forecasting over horizons ranging from seconds to minutes has been addressed only in a limited number of studies (see [22] for a comprehensive review). Most of these approaches rely on local measurements available at these relatively short time scales, such as ground-based irradiance sensors [90], PV power output measurements [63], or sky imagers [3]. Another distinction concerns the choice of the forecast quantity. Common approaches include direct forecasting of PV power generation [67], forecasting solar irradiance [87], and forecasting proxy variables such as the cloud index [8, 35, 84]. Forecasting approaches also differ in whether and how spatial information is incorporated. Many studies consider each PV installation independently and formulate the problem as a purely temporal forecast task. However, irradiance fluctuations induced by cloud motion often exhibit strong spatial coherence and propagate across spatially distributed PV systems [1, 19, 76]. Finally, forecasting approaches differ in whether they predict only the next time step or multiple future time steps. In particular, predictive control requires multi-step forecasts in order to enable proactive control. Multi-step forecasting has therefore received increasing attention in the PV power generation forecasting literature [37, 49, 54].

2-1-2 Classification of forecasting methods

Methods for forecasting PV power generation are commonly classified by their modeling approach [71] into purely physical models, data-driven methods, and mixed approaches. Among purely physical models, numerical weather prediction (NWP) systems provide high-fidelity descriptions of the atmospheric processes that govern solar irradiance. However, in practical operation, the relatively low temporal resolution of NWP outputs and the associated computational effort limit their use for very short term solar forecasting [2, 22]. Consequently, physical models applied at these time scales typically rely on simplified formulations that balance physical detail with computational efficiency [7]. These formulations focus on the dominant physical mechanisms that influence PV power generation, including cloud motion [10, 40], as well as the electrical and thermal behavior of PV panels [6, 26].

In contrast, data-driven methods are constructed from historical data [23]. This class includes time-series models [33] as well as machine learning (ML) models [53]. Such ML models can learn complex nonlinear relationships directly from data and are particularly useful when the underlying physical processes are not fully understood or are expensive to simulate [50]. At

the same time, their performance depends on the availability of large amounts of data at a temporal resolution that is sufficiently high for the intended forecast task [14, 50].

Mixed methods combine elements from different modeling paradigms to exploit their complementary strengths [50]. A particularly promising class of such approaches is physics-informed machine learning (PIML) [42]. For example, [67] incorporates physics-informed features derived from PV performance models into multiple ML architectures. Within the PIML class, PINNs have attracted substantial attention because they incorporate physical equations directly into the training of neural networks (NNs) [68]. Examples include the incorporation of cloud motion dynamics for spatiotemporal forecasting [76], and the modeling of PV panel behavior [88]. The following section reviews the core concepts, recent developments, and key limitations of PINNs.

2-2 Physics-informed neural networks (PINNs)

The class of PIML methods refers to a modeling approach that integrates physical principles with data-driven ML models [42]. Compared to purely data-driven ML models, PIML improves the physical plausibility and interpretability of the resulting forecasts and therefore generalizes better to unseen conditions, such as new spatial locations, forecast horizons, or input configurations. These properties make PIML particularly well suited to settings in which data are sparse, noisy, or limited, where conventional ML models frequently struggle [21]. At the same time, PIML provides advantages over purely physical models when the underlying physical processes are only partially understood or when classical numerical methods face practical limitations. Such limitations may arise from incomplete initial or boundary conditions, uncertain model parameters, or high computational costs [34].

Given these advantages, PIML has been successfully applied in a wide range of scientific and engineering domains, including fluid mechanics [30, 75], materials science [92], climate modeling [43], and power systems [38, 46]. Comprehensive reviews of PIML architectures and applications are available in [24, 42].

Due to the universal function approximation capability of NNs, PINNs have become one of the most studied approaches within the PIML framework. PINNs incorporate physical equations directly into the training of NNs [68]. This enables the model to approximate underlying physical states that not only fit observational data but also adhere to the underlying physical laws. The remainder of this section presents the formal formulation of the PINN framework and discusses existing approaches for collocation point sampling during training.

2-2-1 Formulation of the PINN framework

At their core, PINNs rely on standard feedforward NNs, which consists of an input layer, multiple hidden layers with nonlinear activation functions, and an output layer. Each layer applies a weight matrix and a bias vector to its inputs and maps the result through an activation function. The network parameters, namely the weights and biases, are optimized during training. Following the formulation introduced in one of the foundational works on PINNs [68], the system of interest is commonly described by a nonlinear differential equation

of the form

$$\frac{\partial u(t,\mathbf{x})}{\partial t} = -\mathcal{N}_\gamma[u(t,\mathbf{x})], \quad \mathbf{x} \in \Omega, \quad t \in [0,T], \tag{2-1}$$

where $u(t,\mathbf{x})$ denotes the unknown state function to be approximated, $\mathbf{x}$ is the input vector that may include spatial coordinates and additional variables influencing the state, and γ represents known or unknown system parameters. The operator $\mathcal{N}_\gamma[\cdot]$ is a nonlinear operator derived from the governing physical laws and may involve partial derivatives with respect to $\mathbf{x}$ or t, as well as higher-order derivatives. Ω and $[0,T]$ define the input domain and time interval, respectively.

The state function $u(t,\mathbf{x})$ is approximated by the output of a NN, denoted by $u_\theta(t,\mathbf{x})$, where θ represents the network parameters. To incorporate the governing differential equation into the training process of the PINN, a physics-based residual $\mathcal{R}_\theta(t,\mathbf{x})$ is defined as

$$\mathcal{R}_\theta(t,\mathbf{x}) = \frac{\partial u_\theta(t,\mathbf{x})}{\partial t} + \mathcal{N}_\gamma[u_\theta(t,\mathbf{x})], \tag{2-2}$$

which evaluates to zero when the NN output exactly satisfies the differential equation. The network parameters θ are then optimized by minimizing the overall loss function $\mathcal{L}$, which is defined as the mean squared error (mean squared error (MSE)):

$$\mathcal{L} = \underbrace{\frac{1}{N}\sum_{i=1}^{N}\left(u_\theta(t_i,\mathbf{x}_i) - u(t_i,\mathbf{x}_i)\right)^2}_{\mathcal{L}_{\text{data}}} + \underbrace{\frac{1}{M}\sum_{i=1}^{M}\left(\mathcal{R}_\theta(t_i^{\text{col}},\mathbf{x}_i^{\text{col}})\right)^2}_{\mathcal{L}_{\text{phys}}}. \tag{2-3}$$

Here, $\mathcal{L}_{\text{data}}$ penalizes the discrepancy between the model predictions and the observational data, while $\mathcal{L}_{\text{phys}}$ penalizes violations of the governing differential equation at a set of collocation points. N and M denote the number of observed data points and collocation points, respectively. The collocation points $\{(t_i^{\text{col}},\mathbf{x}_i^{\text{col}})\}_{i=1}^{M}$ are sampled throughout the input domain, while the training data $\{(t_i,\mathbf{x}_i,u(t_i,\mathbf{x}_i))\}_{i=1}^{N}$ typically correspond to initial or boundary conditions. Together, these two loss terms guide the network toward solutions that both fit the observational data and adhere to the underlying physical laws.

Despite advantages such as improved physical plausibility and generalization compared to purely data-driven NNs, PINNs face several practical challenges. These include optimization difficulties caused by complex and competing loss terms, limited performance for multiscale problems or systems with fast dynamics (e.g., stiff equations), and high sensitivity to the choice of network architecture and training hyperparameters [45, 91].

2-2-2 Collocation point sampling

The training of PINN includes a physics-based loss that is evaluated at a set of collocation points sampled from the input domain. The placement of these collocation points influences the convergence behavior and computational cost of the training process, as well as the performance of the resulting forecasts. Early work on PINNs commonly used uniform or random sampling strategies, such as Latin hypercube sampling [68, 89]. Although these strategies are straightforward to implement and provide broad coverage of the domain, they do not account for the varying difficulty of adhering to the governing physical constraints

across different regions of the domain. As a result, they can lead to suboptimal forecast performance and inefficient use of computational resources during training [58]. This has motivated the development of adaptive sampling strategies that aim to place collocation points where they are most beneficial for learning.

Residual-based adaptive sampling methods address this limitation by concentrating collocation points in regions where the physics residual is large, indicating that the model adheres poorly to the governing physical constraints in these regions [89]. More advanced adaptive strategies further refine the placement of collocation points using derivatives of the residual [18, 31]. Closely related importance sampling approaches bias the sampling distribution toward regions that contribute most to the loss function [58].

Despite the wide range of adaptive sampling strategies, approaches that account for the statistical structure of the input space remain largely unexplored. Ignoring dependencies between input variables can produce physically implausible collocation points and degrade both training efficiency and overall model performance.

2-3 Multi-step and recursive forecasting

Multi-step forecasting is a fundamental problem in time-series analysis and refers to the forecasting of a sequence of future values from a given reference time [52]. The difficulty arises from the increasing uncertainty of forecasts over the forecast horizon and from the propagation of forecast errors over successive forecast steps [12, 82].

Several strategies have been developed for multi-step forecasting. Each strategy involves trade-offs in terms of model complexity and computational cost. The recursive strategy [82] trains a single model to forecast one step into the future. During inference, this model is applied iteratively to generate multi-step forecasts, with previously generated forecasts fed back as inputs. Although this approach is computationally efficient, it is susceptible to error accumulation since errors introduced at early steps propagate and amplify over the forecast horizon [12, 85].

In contrast, the direct strategy [44] trains a separate model for each forecast step of interest. This approach avoids the iterative propagation of errors inherent in recursive forecasting, as each step is forecast independently. However, it requires training and maintaining multiple models, which increases both the training complexity and computational burden, particularly for long forecast horizons with fine temporal resolution. Furthermore, the direct strategy does not explicitly model the dependencies between successive steps, which may limit the temporal coherence of multi-step forecasts [82].

Hybrid strategies that combine elements of both recursive and direct forecasting have also been proposed. These include strategies that combine direct and recursive forecasting, in which separate models are used for different segments of the forecast horizon and recursive forecasting is applied within each segment [81]. In addition, approaches with multiple outputs train a single model to forecast several future steps simultaneously [15].

The choice of multi-step forecasting strategy is influenced by factors including the forecast horizon, the temporal resolution of available data and the resolution required by the intended application, computational constraints, and the characteristics of the underlying system being

modeled. Despite its susceptibility to error accumulation, the recursive strategy remains widely used because of its simplicity and computational efficiency.

Recurrent NNs and related sequential architectures are the dominant modeling framework for recursive inference, as they are specifically designed to capture temporal dependencies in sequential data. However, in standard practice these models are typically trained using one-step-ahead prediction with teacher forcing, creating a mismatch between training and recursive inference. To mitigate error accumulation under recursive inference, multi-step training objectives have been introduced in which the models are trained to minimize their own recursive multi-step forecast errors, such as scheduled sampling [13].

In addition, recent research has explored the incorporation of physical knowledge into data-driven methods to improve multi-step forecasting under recursive inference. For example, the work in [65] introduces physical constraints into neural time-series models, while the work in [59] reformulates PINNs within an autoregressive framework. These approaches illustrate the growing interest in leveraging physical principles to enhance the stability of multi-step inference.

Despite these advances, most existing work on recursive forecasting relies on dense temporal supervision during training. In many practical applications, however, multi-step forecasts are required at a higher temporal resolution than is available in the training data. Consequently, recursive forecasting under temporal data sparsity remains an open research problem.

2-4 Summary

This chapter reviewed the background and literature on PV power generation forecasting for real-time power system control. It first discussed the operational requirements of modern power systems and the key design choices involved in defining the forecasting task. The main classes of existing forecasting methods were presented, including physical, data-driven, and mixed approaches. Next, the chapter introduced PINNs and discussed existing approaches for collocation point sampling. Finally, it reviewed the main strategies for multi-step forecasting and highlighted the fundamental challenges of recursive forecasting, including error accumulation and the reliance of existing methods on dense temporal supervision. These limitations motivate the development of improved recursive forecasting frameworks that can operate under temporal data sparsity, which is the central focus of this thesis.

Chapter 3

Methodology of the PINN-based forecasting framework

This chapter presents the methodological framework developed to investigate the use of physics-informed neural networks (PINNs) to forecast photovoltaic (PV) power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. Specifically, the forecasting task is formulated as a recursive multi-step high-resolution prediction problem that bridges this temporal resolution mismatch.

The remainder of this chapter is structured as follows. Section 3-1 introduces the PINN formulation, including the governing equation used to encode partial physical knowledge of atmospheric processes relevant to PV power generation, specifically spatiotemporal cloud motion. Section 3-2 presents a correlation-based collocation point sampling strategy developed in this work. Section 3-3 describes the normalization of input variables and the scaling of the governing cloud motion equation. Section 3-4 proposes a recursive forecasting framework to bridge temporal data sparsity, while Section 3-5 proposes an improved training strategy to align training with recursive inference. Section 3-6 describes the baseline models used for comparison, and Section 3-7 details the validation setup.

3-1 PINN design

This section describes the formulation of the PINN used for PV power generation forecasting in spatially distributed PV systems. The design focuses on three aspects: the definition of the forecast target and model inputs, the incorporation of a governing spatiotemporal cloud motion equation, and the construction of a composite training loss function that balances agreement with observational data and physics-based regularization.

3-1-1 Forecast target and model inputs

At intra-hour temporal scales, variability in PV power generation is largely driven by spatiotemporal cloud motion and the resulting modulation of solar irradiance [22]. Therefore, the PINN is designed to predict the cloud index, which provides a physically meaningful representation of irradiance variability and remains directly applicable for downstream PV power generation estimation.

The cloud index c is defined as

$$c = \frac{I}{I_{\text{clearsky}}}, \tag{3-1}$$

where I denotes the global horizontal irradiance (GHI) and I_{clearsky} denotes the corresponding clear-sky GHI. GHI is chosen as the basis for the cloud index formulation because it quantifies the total solar irradiance incident on a horizontal surface, combining both direct and diffuse components, and is therefore directly relevant to PV power generation. Normalization by the clear-sky irradiance removes diurnal and seasonal effects associated with solar geometry, thereby isolating irradiance variability attributable to cloud cover [8, 35, 84].

The cloud index is bounded, with $c \in [0, 1]$, where values close to one indicate clear-sky conditions and lower values indicate increased cloud attenuation. After forecasting, irradiance can be reconstructed by multiplying the predicted cloud index with I_{clearsky}, which can be computed from solar geometry based on time and location [41].

The PINN represents a neural network parameterized by θ whose output defines the function

$$c_\theta : \mathcal{Q} \to \mathbb{R}, \tag{3-2}$$

which approximates the underlying cloud index function c, where $\mathcal{Q}$ denotes the domain of the input variables. The input vector q is defined as

$$\mathbf{q} = \left[x,\, y,\, t,\, T_{\text{air}},\, Z,\, u,\, v,\, \{c_l\}_{l=1}^{L}\right], \tag{3-3}$$

where (x, y) denote the spatial coordinates of the target location and t denotes the forecast horizon (i.e., the forecast lead time). The air temperature T_{air} provides contextual information on atmospheric conditions and has been shown to be a relevant auxiliary input variable for PV power generation forecasting [32]. The solar zenith angle Z encodes the position of the sun relative to the horizon and captures systematic diurnal and seasonal variations in irradiance geometry. The wind velocity components (u, v) govern cloud motion and enter directly into the cloud motion equation described in Section 3-1-2. The set $\{c_l\}_{l=1}^{L}$ denotes cloud index observations at L fixed measurement locations within the spatial domain under consideration. Together, these observations characterize the spatial distribution of cloud cover at the reference time and provide information on upstream cloud patterns that are relevant for forecasting the cloud index at the target location.

All input variables are observed at the reference time t_0, while the model output, corresponding to the target cloud index c_*, is defined at the target time $t_* = t_0 + t$. An overview of all variables, including their temporal and spatial interpretation and units, is provided in Table 3-1.

Table 3-1: Overview of input and output variables. The column "Time" indicates whether a variable is observed at the reference time t_0 or corresponds to the target time t_*. The column "Location" indicates the spatial position at which the variable is defined, either the target location (x, y) or the L measurement locations.

	Description	**Time**	**Location**	**Unit**
Inputs				
(x, y)	Target location	-	-	m
T_{air}	Air temperature	t_0	(x, y)	°C
Z	Solar zenith angle	t_0	(x, y)	°
(u, v)	Wind components	t_0	(x, y)	m/s
$\{c_l\}_{l=1}^{L}$	Cloud index observations	t_0	L sites	-
t	Forecast horizon	-	-	s
Output				
c_*	Target cloud index	t_*	(x, y)	-

3-1-2 Physical model: cloud motion

Partial physical knowledge of the atmospheric processes governing irradiance variability is incorporated through a spatiotemporal cloud motion model. Specifically, the cloud index c is assumed to satisfy a cloud advection partial differential equation of the form [56, 69]:

$$\frac{\partial c}{\partial t} + u\frac{\partial c}{\partial x} + v\frac{\partial c}{\partial y} = 0. \tag{3-4}$$

This equation describes the transport of the cloud index through space and time driven by the wind velocity components (u, v) and captures advection-dominated cloud motion. This assumption is appropriate for the temporal scales considered in this thesis, which are on the order of minutes, because internal cloud evolution processes are comparatively slow over such horizons [22]. Within the PINN framework, the cloud motion equation is incorporated during training as a physics-based regularization term that penalizes violations of this equation.

3-1-3 Training loss function

Training the PINN consists of minimizing a composite loss function that balances agreement with observational data and adherence to the governing cloud motion equation. The total training loss is defined as

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}}\,\mathcal{L}_{\text{phys}}, \tag{3-5}$$

where $\lambda_{\text{phys}} > 0$ is a weighting parameter that controls the strength of the physics-based regularization.

The data loss penalizes discrepancies between the network predictions $c_\theta(\mathbf{q})$ and the corresponding observed target values c_* using the mean squared error (MSE) and is defined as

$$\mathcal{L}_{\text{data}} = \frac{1}{N}\sum_{i=1}^{N}\left(c_\theta(\mathbf{q}_i) - c_{*,i}\right)^2, \tag{3-6}$$

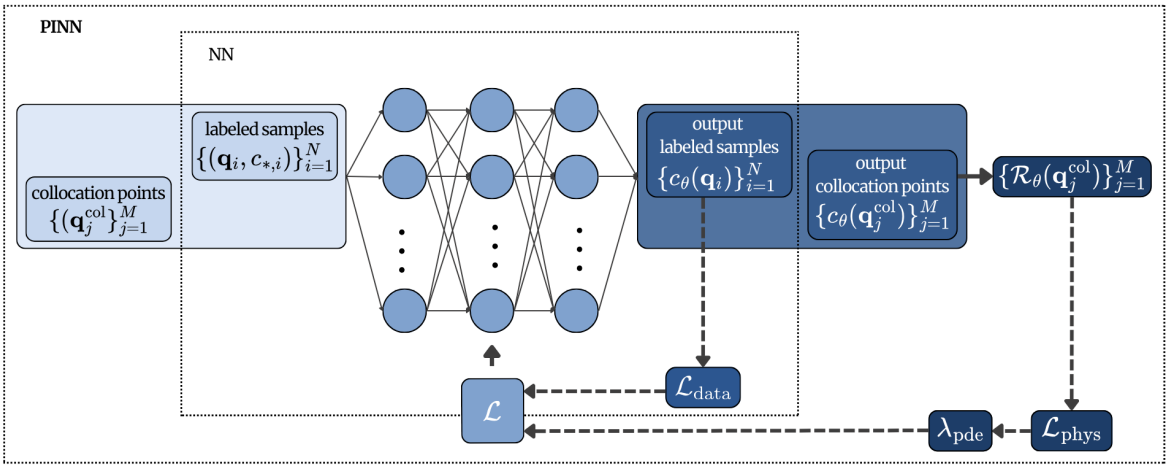

Figure 3-1: Overview of the PINN training framework. A neural network (NN) learns an approximate solution of the cloud index function c_θ by minimizing a composite loss $\mathcal{L}$. The loss consists of a data term $\mathcal{L}_{\text{data}}$ computed from labeled training samples $\{(\mathbf{q}_i, c_{*,i})\}_{i=1}^{N}$, and a physics term $\mathcal{L}_{\text{phys}}$ computed from collocation points $\{\mathbf{q}_j^{\text{col}}\}_{j=1}^{M}$ through the residual $\mathcal{R}_\theta$ of the governing cloud motion equation.

where $\{(\mathbf{q}_i, c_{*,i})\}_{i=1}^{N}$ denotes the set of N labeled training samples.

The physics loss penalizes violations of the cloud motion equation at a set of collocation points by using the mean squared residual of the governing equation evaluated on the network predictions, and is defined as

$$\mathcal{L}_{\text{phys}} = \frac{1}{M}\sum_{j=1}^{M}\left(\mathcal{R}_\theta(\mathbf{q}_j^{\text{col}})\right)^2, \tag{3-7}$$

where $\{\mathbf{q}_j^{\text{col}}\}_{j=1}^{M}$ denotes the set of M collocation points. These points are synthetically generated input samples that span the input domain. Unlike labeled training samples, collocation points do not require corresponding target values and are used only to enforce adherence to the governing equation during training. The developed procedure for generating collocation points is described in detail in Section 3-2.

The residual of the governing cloud motion equation, based on Equation 3-4, is defined as

$$\mathcal{R}_\theta(\mathbf{q}^{\text{col}}) = \frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial t} + u\frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial x} + v\frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial y}. \tag{3-8}$$

The explicit form of the residual depends on the normalization of the input variables, which is addressed in Section 3-3. All partial derivatives are computed using automatic differentiation [11].

An overview of the PINN training framework, including the flow of labeled training samples and collocation points through the network and the computation of the total loss, is illustrated in Figure 3-1.

3-2 Collocation point sampling strategy

Collocation points play a central role in PINNs, as they determine the locations in the input space at which the governing equations are enforced. In this thesis, collocation points are used to regularize the learned cloud index dynamics by enforcing the cloud motion equation described in Section 3-1-2 during training. Their construction is therefore an important design choice that directly influences both the physical plausibility of the resulting forecasts and the generalization performance of the PINN.

Existing methods often rely on independent sampling of each variable [89]. However, a fundamental challenge in generating collocation points arises when the input space comprises multiple variables that are both statistically dependent and physically coupled through the underlying atmospheric processes, as in PV power generation forecasting. In such settings, independent sampling can produce combinations that are physically implausible and do not occur in reality. In contrast, fully joint sampling based on data is often poorly representative of the true underlying distribution in regimes with relatively little data, particularly as dimensionality increases [73]. Both limitations can lead to misleading physics-based regularization and inefficient use of model capacity. These considerations motivate the development of intermediate sampling strategies that simultaneously promote physical plausibility and statistical representativeness.

To address this trade-off, this thesis proposes a novel collocation point sampling strategy that exploits statistical correlations among the input variables. The approach preserves the main dependency structures by grouping variables into coherent clusters based on their statistical dependence. Collocation points are then generated by sampling variables jointly within each cluster, while sampling different clusters independently. This reduces the effective dimensionality of the joint sampling problem, thereby mitigating both the generation of physically implausible input combinations and the poor statistical representativeness associated with high-dimensional joint sampling. More broadly, the proposed sampling strategy is applicable to PINN formulations with statistically dependent inputs, independent of the specific physical system or governing equations under consideration.

The developed sampling procedure consists of four steps: correlation analysis, input variable clustering, cluster-wise joint sampling, and temporal and spatial sampling. Each step is described in detail below.

3-2-1 Correlation analysis

The sampling procedure begins with an analysis of statistical dependencies among the input variables. Dependencies are quantified using a Pearson correlation matrix [47] computed from the observed training data. Pearson correlation is chosen for its simplicity and its ability to capture dominant linear relationships among the input variables.

Let $\mathbf{Q} \in \mathbb{R}^{n \times d}$ denote the data matrix constructed from the training data, where n is the number of training samples and d is the number of input variables included in the correlation analysis. In the context of this thesis, the analysis includes meteorological variables, whereas spatial coordinates (x, y) and the forecast horizon t are sampled independently. Each column $\mathbf{q}^j \in \mathbb{R}^n$ corresponds to one input variable, $j \in \{1, \ldots, d\}$.

The Pearson correlation coefficient between variables j and m is defined as [47]

$$\rho_{jm} = \frac{\text{cov}(\mathbf{q}^j, \mathbf{q}^m)}{\sigma_{\mathbf{q}^j}\, \sigma_{\mathbf{q}^m}}, \tag{3-9}$$

where $\text{cov}(\cdot, \cdot)$ denotes the sample covariance and $\sigma_{\mathbf{q}^j}$ is the sample standard deviation of variable j. The resulting correlation matrix provides a basis for identifying groups of related variables.

3-2-2 Correlation-based input variable clustering

To transform the correlation matrix into a form suitable for clustering, it is converted into a dissimilarity matrix

$$D_{jm} = 1 - |\rho_{jm}|, \tag{3-10}$$

where strongly correlated variables (either positively or negatively) have low dissimilarity. This transformation reflects the objective of grouping variables based on the strength of their statistical dependence, regardless of the sign of the correlation.

Hierarchical agglomerative clustering [51, 60] is then applied to the dissimilarity matrix. This bottom-up procedure initially treats each variable as a separate cluster and iteratively merges clusters based on their average pairwise dissimilarity. This produces a hierarchical clustering structure represented by a dendrogram, which visualizes the pairwise similarities among the input variables. The average linkage criterion used to merge clusters A and B is defined as [60]

$$L(A, B) = \frac{1}{|A||B|} \sum_{j \in A} \sum_{m \in B} D_{jm}, \tag{3-11}$$

where $|A|$ and $|B|$ denote the number of variables in each cluster.

A dendrogram can be cut at different levels to obtain the desired number of clusters k. In this work, the dendrogram is cut at an empirically chosen level to obtain clusters that correspond to distinct and physically interpretable processes relevant to PV power generation forecasting. These processes include cloud-related variability, atmospheric thermodynamic effects, wind-driven transport, and solar geometry.

3-2-3 Cluster-wise joint sampling

Once the input variable clusters have been identified, collocation points are generated by sampling each cluster independently, while sampling all variables within a cluster jointly. This design reflects the assumption that strong dependencies exist within clusters, while dependencies between clusters are weaker and can be neglected for the purpose of collocation point sampling.

Sampling is performed separately for each calendar month to account for seasonal variability and to reduce the occurrence of unrealistic combinations of variables arising from independently sampled clusters. Within each month, joint samples for each cluster are drawn using multivariate kernel density estimation (KDE).

Given training samples $\{\mathbf{q}_i\}_{i=1}^{n} \subset \mathbb{R}^{d_c}$ for a cluster of dimension d_c, the multivariate KDE is defined as [20]

$$\hat{f}(\mathbf{q}) = \frac{1}{nh^{d_c}} \sum_{i=1}^{n} K\left(\frac{\mathbf{q} - \mathbf{q}_i}{h}\right), \tag{3-12}$$

where $\mathbf{q} \in \mathbb{R}^{d_c}$ denotes a generic evaluation point, $h > 0$ is a bandwidth parameter that controls the smoothness of the estimated density, and $K(\cdot)$ is a kernel function. In this work, a multivariate Gaussian kernel is used:

$$K(\mathbf{z}) = (2\pi)^{-d_c/2} \exp\left(-\frac{1}{2}\|\mathbf{z}\|^2\right), \tag{3-13}$$

where $\mathbf{z} \in \mathbb{R}^{d_c}$.

A KDE provides a flexible, nonparametric approximation of the joint probability density function, implicitly capturing the empirical marginal distributions of the variables as well as their joint dependency structure. Although the performance of KDE degrades with increasing dimensionality due to the curse of dimensionality [73], empirical results in this work indicate that KDE-based sampling remains sufficiently accurate for the cluster sizes considered. More complex dependence modeling approaches were also evaluated, but did not yield consistent performance improvements and are therefore omitted for clarity.

After cluster-wise sampling, the sampled variables are concatenated to form complete synthetic meteorological input vectors. Samples drawn from the tails of the KDE that violate basic physical constraints are discarded. In the context of this thesis, this includes restricting the solar zenith angle to daylight conditions,

$$0^\circ \leq Z < 90^\circ,$$

and enforcing physically meaningful bounds on the cloud index,

$$0 < c_l \leq 1 \quad \forall l \in \{1, \ldots, L\}.$$

3-2-4 Temporal and spatial sampling

Finally, each synthetic meteorological input sample is augmented with spatial coordinates (x, y) and a forecast horizon t. These quantities are sampled independently of the meteorological variables.

Forecast horizons are drawn from a uniform distribution

$$t \sim \mathcal{U}(0, t_{\max} + \delta_t), \tag{3-14}$$

where $t_{\max}$ corresponds to the maximum forecast horizon of interest, and δ_t denotes the small additional margin to extend the temporal domain over which the governing equations are enforced.

Spatial coordinates are sampled uniformly over a two-dimensional region defined by the convex hull of the L measurement locations, expanded by a spatial margin δ_s:

$$(x, y) \sim \mathcal{U}(\mathcal{H}_{\delta_s}), \tag{3-15}$$

where $\mathcal{H}_{\delta_s}$ denotes the buffered convex hull. Sampling over this buffered region allows the governing equations to be enforced slightly beyond the spatial extent of the observed training samples, supporting spatial generalization of the PINN.

Overall, the developed sampling procedure provides a broadly applicable, correlation-based approach for generating collocation points. It generates physically plausible and statistically representative collocation points, thereby supporting effective physics-based regularization during PINN training.

3-3 Normalization and scaling of inputs and governing equation

Normalization of input variables plays a critical role in training and numerical stability of PINNs [17], serving three closely related purposes. First, it improves numerical conditioning and optimization stability by preventing input variables with large numerical scales from dominating gradient magnitudes during training. Second, it promotes a balanced contribution of the different terms in the residual of the governing cloud motion equation, which involves partial derivatives with respect to multiple input variables. Third, it aligns the numerical scale of the physics loss with that of the data loss, enabling interpretable and stable tuning of the physics-loss weighting parameter λ_{phys}.

As a consequence, normalization of the input variables directly affects the scale of the associated partial derivatives. The governing cloud motion equation must therefore be reformulated consistently in the normalized input space to ensure a numerically well-conditioned residual that remains consistent with the underlying physical model.

3-3-1 Input normalization

All input variables, except the cloud index observations $\{c_l\}_{l=1}^{L}$, which are already dimensionless and bounded in $[0, 1]$, are normalized to the unit interval prior to model training using min-max scaling (i.e., each variable is rescaled by its observed range). The normalization bounds are determined from the available training data for the spatiotemporal domain under consideration. This normalization is applied because the input variables differ in physical units and numerical ranges, and min-max scaling maps these heterogeneous inputs to a common bounded interval without imposing assumptions on their underlying distributions. The model output is not normalized because the cloud index is already expressed in the unit interval.

For a generic input variable q with bounds $q_{\min}$ and $q_{\max}$, the normalized variable q' is defined as

$$q' = \frac{q - q_{\min}}{\Delta q}, \tag{3-16}$$

where $\Delta q = q_{\max} - q_{\min}$. Both the observed training samples and the collocation points are normalized using this transformation, resulting in scaled inputs $\mathbf{q}'$ and $\mathbf{q}^{\text{col}'}$, respectively. The same normalization parameters are applied during inference, ensuring that all inputs are mapped to the same normalized space as during training.

3-3-2 Governing equation scaling

Accordingly, the governing cloud motion equation described in Section 3-1-2 is reformulated in the normalized input space. Applying the chain rule, partial derivatives with respect to the original spatial and temporal variables transform as

$$\frac{\partial}{\partial x} = \frac{1}{\Delta x}\frac{\partial}{\partial x'}, \quad \frac{\partial}{\partial y} = \frac{1}{\Delta y}\frac{\partial}{\partial y'}, \quad \frac{\partial}{\partial t} = \frac{1}{\Delta t}\frac{\partial}{\partial t'}. \tag{3-17}$$

The physical wind components are recovered from their normalized counterparts via the inverse min-max transformation

$$u = u_{\min} + u'\,\Delta u, \quad v = v_{\min} + v'\,\Delta v. \tag{3-18}$$

Substituting the transformed partial derivatives and recovered wind components into the original cloud motion equation yields the following scaled formulation:

$$\frac{1}{\Delta t}\frac{\partial c}{\partial t'} + (u_{\min} + u'\,\Delta u)\,\frac{1}{\Delta x}\frac{\partial c}{\partial x'} + (v_{\min} + v'\,\Delta v)\,\frac{1}{\Delta y}\frac{\partial c}{\partial y'} = 0. \tag{3-19}$$

To obtain a nondimensional form consistent with the data loss, the equation is rescaled by multiplying by Δt, yielding

$$\frac{\partial c}{\partial t'} + \tilde{u}\frac{\partial c}{\partial x'} + \tilde{v}\frac{\partial c}{\partial y'} = 0, \tag{3-20}$$

where the dimensionless velocity coefficients are defined as

$$\tilde{u} = \frac{u\,\Delta t}{\Delta x}, \quad \tilde{v} = \frac{v\,\Delta t}{\Delta y}. \tag{3-21}$$

This transformation yields a numerically balanced residual that is more suitable for gradient-based optimization and operates on a scale comparable to the data loss. The final residual used in the loss function is therefore given by

$$\mathcal{R}_\theta(\mathbf{q}^{\mathrm{col}'}) = \frac{\partial c_\theta(\mathbf{q}^{\mathrm{col}'})}{\partial t'} + \tilde{u}\frac{\partial c_\theta(\mathbf{q}^{\mathrm{col}'})}{\partial x'} + \tilde{v}\frac{\partial c_\theta(\mathbf{q}^{\mathrm{col}'})}{\partial y'}. \tag{3-22}$$

3-4 Recursive forecasting framework

This section proposes a recursive forecasting framework designed to address the mismatch between the temporally sparse available data and the higher-resolution multi-step forecasts required in power system control applications.

3-4-1 Forecasting procedure during inference

Consider a PINN that is trained on temporally sparse observations with temporal resolution t_{train} and has output c_θ. With a slight abuse of notation, the information available about the system at a reference time t_0 is represented by a set of L input vectors

$$\{\mathbf{q}_\ell(t_0)\}_{\ell=1}^{L} \in \mathbb{R}^{d+L},$$

corresponding to the L measurement locations in the spatial domain under consideration. Each vector contains a set of d local auxiliary variables $\mathbf{m}_\ell(t_0)$, which include spatial coordinates, the forecast horizon, and meteorological variables. These variables are assumed to remain constant or exogenous over the forecast horizon. In addition, each vector includes the cloud index observations from all L locations $\{c_j(t_0)\}_{j=1}^{L}$ that evolve over the forecast horizon and together form a spatial cloud index field

$$\mathbf{c}(t_0) = \{c_j(t_0)\}_{j=1}^{L}.$$

Each input vector is then given by

$$\mathbf{q}_\ell(t_0) = [\mathbf{m}_\ell(t_0),\, \mathbf{c}(t_0)], \quad \ell = 1, \ldots, L.$$

Importantly, the cloud index field $\mathbf{c}(t_0)$ is identical across all input vectors and therefore constitutes a shared set of state variables. Updating this shared state requires the simultaneous generation of cloud index forecasts at all locations, and recursive forecasting is therefore performed jointly across all input vectors.

To generate multi-step forecasts at a temporal resolution higher than that of the training data, the trained PINN is applied recursively over a forecast horizon t_r, corresponding to the desired inference-time temporal resolution. In the first recursive step, the model is evaluated on all input vectors to generate a set of cloud index forecasts, which collectively form an updated estimate of the cloud index field:

$$\hat{\mathbf{c}}(t_0 + t_r) = \{c_\theta(\mathbf{q}_\ell(t_0))\}_{\ell=1}^{L}. \tag{3-23}$$

This is fed back into all input vectors to replace the state variables and construct new inputs

$$\mathbf{q}_\ell(t_0 + t_r) = [\mathbf{m}_\ell(t_0),\, \hat{\mathbf{c}}(t_0 + t_r)], \quad \ell = 1, \ldots, L.$$

This recursive update procedure is repeated for a desired number of steps $k \in \{1, \ldots, N_r\}$:

$$\hat{\mathbf{c}}(t_0 + kt_r) = \{c_\theta(\mathbf{q}_\ell(t_0 + (k-1)t_r))\}_{\ell=1}^{L}\,, \tag{3-24}$$

thereby producing a multi-step forecast of length N_r at a temporal resolution t_r.

Throughout this process, no new observations are assumed to be available beyond the input at the reference time t_0, due to temporal data sparsity. As a result of recursive updating without intermediate supervision, forecast errors can accumulate over successive steps, making multi-step forecast performance sensitive to the model's ability to generate temporally coherent forecasts under recursive use.

This recursive forecasting procedure can, in principle, be applied to a wide range of learning-based forecasting models. However, because PINNs incorporate governing equations in the training loss, they can generalize between sparse observations and produce physically plausible multi-step forecasts under recursive deployment. This makes PINNs particularly well suited for use within the developed recursive forecasting framework.

3-4-2 Forecasting procedure for validation

Validating multi-step forecasts at a temporal resolution higher than that of the available data is an inherent limitation in practical applications, as forecast errors cannot be evaluated at intermediate recursive steps. To address this limitation, validation is performed by executing the recursive forecasting procedure over the coarse interval t_{train} supported by the available data, corresponding to $N_r = t_{\text{train}}/t_r$ recursive steps, and comparing the terminal state forecast

$$\hat{\mathbf{c}}(t_0 + t_{\text{train}}),$$

with the corresponding observed target at time $t_0 + t_{\text{train}}$

$$\mathbf{c}(t_0 + t_{\text{train}}),$$

which is available in the labeled dataset.

This terminal forecast comparison assesses error accumulation over the coarse interval of the available data and provides a measure of the model's performance under recursive deployment at higher temporal resolution. This approach implicitly assumes that producing an accurate terminal forecast is a meaningful proxy for the quality of the underlying recursively generated forecast trajectory. In the absence of intermediate labeled observations, this assumption is not only reasonable but practically one of the few methods to evaluate the model.

3-5 Improved training strategy using recursive forecasting

To address the mismatch between the temporally sparse available data and the higher-resolution multi-step forecasts required in control applications, models trained on temporally sparse data are commonly deployed recursively at higher temporal resolution during inference, as exemplified by the recursive forecasting framework developed in Section 3-4.

This discrepancy between different temporal resolutions introduces a training-inference mismatch: while the model is optimized for single-pass prediction accuracy at the temporal resolution of the available data, its inference-time higher-resolution multi-step performance depends on its behavior under repeated recursive application. This section addresses this mismatch by proposing an improved training strategy that explicitly incorporates the recursive forecasting framework into the training process, thereby aligning training with recursive inference.

3-5-1 Training procedure

Consider again a PINN with output c_θ and training data consisting of N labeled samples available at a temporally sparse resolution t_{train}. Each labeled sample i corresponds to a reference time t_i and consists of a set of L input vectors

$$\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L} = \{[\mathbf{m}_\ell(t_i),\, \mathbf{c}(t_i)]\}_{\ell=1}^{L},$$

where $\mathbf{m}_\ell$ and $\mathbf{c}$ are defined in Section 3-4, together with the corresponding observed target cloud index field

$$\mathbf{c}_{*,i} = \mathbf{c}(t_i + t_{\text{train}}).$$

During training, the PINN is applied recursively to each training sample over $N_r = t_{\text{train}}/t_r$ steps, following the recursive forecasting procedure described in Section 3-4. Cloud index forecasts are generated jointly across all input vectors, which defines an N_r-step joint recursive rollout operator

$$\mathcal{C}_\theta^{(N_r)}(\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L}) = \hat{\mathbf{c}}(t_i + N_r\, t_r),$$

mapping a set of input vectors to the terminal forecast of the cloud index field, where $\hat{\mathbf{c}}$ is defined in Equation 3-23.

The recursive data loss is defined by the MSE between the terminal recursive forecasts and the corresponding observed cloud index fields:

$$\mathcal{L}_{\text{data}}^{\text{rec}} = \frac{1}{NL}\sum_{i=1}^{N}\left\|\mathcal{C}_\theta^{(N_r)}\left(\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L}\right) - \mathbf{c}_{*,i}\right\|_2^2. \tag{3-25}$$

By comparing only the terminal forecasts, this loss implicitly captures the accumulation of errors across recursive steps. As a result, the training objective is aligned with inference-time deployment, encouraging stable and accurate performance under recursive use at higher temporal resolution. The physics loss remains unchanged and is computed independently using collocation points, as described in Section 3-1-3.

3-5-2 Discussion

The improved training strategy proposed in this section is applicable to a broad class of learning-based forecasting models. It enables training to be aligned with recursive inference in settings where available data are temporally sparse and higher-resolution multi-step forecasts are required.

A practical limitation of this training strategy is the increased computational cost associated with backpropagation through multiple recursive steps. Consequently, the choice of the recursive forecast horizon t_r used during training introduces a trade-off between computational cost and alignment with inference-time deployment. Increasing t_r reduces the number of recursive steps $N_r = t_{\text{train}}/t_r$ and therefore reduces computational training cost. However, this may insufficiently expose the model to the higher-resolution rollouts required during inference, limiting its ability to mitigate error accumulation during deployment. In contrast, choosing t_r to match the temporal resolution required during inference yields full alignment between training and deployment, but at the expense of increased computational cost.

3-6 Comparison models

Two PINN formulations are proposed in this thesis:

- *Baseline PINN*: trained using conventional single-pass supervision, as described in Section 3-1-3;
- *Improved PINN*: trained using the improved strategy based on recursive forecasting, proposed in Section 3-5.

To evaluate the forecast performance of these PINN formulations for control-relevant PV power generation forecasting under temporal data sparsity, three additional comparison models are considered: a persistence model, purely data-driven NNs, and a physics-only cloud motion model. For consistency with the two PINN formulations, two corresponding NN variants are used: a *baseline NN* trained with single-pass supervision, and an *improved NN* trained using the improved strategy. All comparison models are described in the following sections.

3-6-1 Persistence model

The persistence model is included as a reference model that assumes the forecast variables remain constant over the prediction horizon, against which the added value of dynamical forecasting models can be evaluated. Formally, the cloud index at all future times is set equal to the observed cloud index at the reference time,

$$\hat{\mathbf{c}}(t_0 + t_r) = \mathbf{c}(t_0). \tag{3-26}$$

Persistence constitutes a strong reference for forecasting over intervals t_{train} on the order of minutes, as cloud fields evolve only relatively slowly on such time scales. As a result, persistence can yield highly competitive terminal forecast accuracy in low-variability regimes, leaving limited room for improvement by more complex models. However, because persistence neglects cloud motion and other atmospheric processes, its performance deteriorates during periods of higher variability. In addition, its performance degrades when the temporal resolution of the available data t_{train} is too coarse.

3-6-2 Data-driven neural network models

Two purely data-driven NN variants are considered: a *baseline NN* and an *improved NN*. They are included to evaluate the predictive capability of data-driven learning under temporal data sparsity. Both models use the same network architecture, input variables, and output definition as the PINN formulations, but are trained without physics-based regularization, i.e., $\mathcal{L} = \mathcal{L}_{\text{data}}$.

Assuming equal training procedures and hyperparameters, any performance differences between models and their corresponding PINN counterparts can be attributed directly to the inclusion of physics-based regularization during training. However, because these NN models are trained on temporally sparse data and lack of constraints that promote temporally coherent forecasts, their recursive deployment at higher temporal resolution can lead to error accumulation.

3-6-3 Physics-only cloud motion model

A physics-only model based on the cloud motion equation described in Section 3-1-2 is included as a comparison model to evaluate the predictive capability of partial physical knowledge alone. Under the assumption of constant wind fields over the forecast horizon, the cloud motion equation admits the closed-form solution [29]:

$$c(x, y, t + t_r) = c(x - u\,t_r,\ y - v\,t_r,\ t). \tag{3-27}$$

This translates the current cloud field forward in time along the wind direction to obtain future fields. In practice, this solution is evaluated numerically using a semi-Lagrangian backtracking approach [83].

By construction, this model enforces physically plausible cloud motion but lacks any data-driven mechanism to compensate for modeling and numerical approximations. Consequently, its recursive deployment can be sensitive to the simplifying assumption of purely advective cloud dynamics, as well as to the interpolation and extrapolation approximations required under sparse spatial observations.

3-7 Validation design

This section defines the validation design used to evaluate the forecast performance of the two proposed PINN formulations for control-relevant PV power generation forecasting under temporal data sparsity. Their forecast performance is compared with the comparison models introduced in Section 3-6. The validation design ensures a fair, control-relevant, and statistically robust evaluation across all models within the recursive forecasting framework developed in Section 3-4.

3-7-1 Reproducibility

To ensure a fair comparison, all models are evaluated under identical experimental conditions, including the use of the same recursive forecasting procedure, identical test datasets, and a common set of evaluation metrics. For learning-based models, identical network architectures and training datasets are used. Together, these experimental design choices ensure that performance differences between models can be attributed to modeling and training choices rather than differences in data or validation design.

To account for stochastic variability, the entire validation pipeline is repeated over 10 independent random seeds. For each seed, all stochastic components in the pipeline are initialized using the same seed value. These components include: (i) the partitioning of the dataset into training and test sets, (ii) the generation of collocation points used in the physics loss, and (iii) NN weight initialization. Each seed therefore corresponds to a fully independent realization of model training and evaluation, while reproducibility within each realization is ensured.

3-7-2 Evaluation metrics

Forecast performance is quantified by comparing terminal recursive cloud index field forecasts $\hat{\mathbf{c}}$ with corresponding observed cloud index field targets $\mathbf{c}_*$. The following error metrics are used:

- **Root mean squared error (RMSE)**, which emphasizes larger forecast errors:

$$\mathrm{MSE} = \sqrt{\frac{1}{NL}\sum_{i=1}^{N} \|\hat{\mathbf{c}}_i - \mathbf{c}_{*,i}\|_2^2}. \tag{3-28}$$

- **Mean absolute error (MAE)**, which captures absolute forecast errors:

$$\mathrm{MAE} = \frac{1}{NL} \sum_{i=1}^{N} \left\| \hat{\mathbf{c}}_i - \mathbf{c}_{*,i} \right\|_1 . \tag{3-29}$$

- **Bias**, which captures systematic overestimation or underestimation of forecasts:

$$\mathrm{Bias} = \frac{1}{NL} \sum_{i=1}^{N} \sum_{l=1}^{L} \left(\hat{c}_{i,l} - c_{*,i,l} \right) . \tag{3-30}$$

- **Variability error (VE)**, which assesses the preservation of realistic spatial variability in forecasts:

$$\mathrm{VE} = \frac{1}{N} \sum_{i=1}^{N} \left| \mathrm{Var}(\hat{\mathbf{c}}_i) - \mathrm{Var}(\mathbf{c}_{*,i}) \right| , \tag{3-31}$$

  where $\mathrm{Var}(\cdot)$ denotes the variance across the L spatial measurement locations within a forecast instance.

These metrics are reported as the mean and standard deviation over the 10 independent training seeds. This aggregation provides a measure of both average performance and robustness to stochastic variability arising from data sampling and model initialization.

3-8 Summary

This chapter presented the methodological framework developed in this thesis to forecast PV power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. The forecasting task was formulated as a multi-step prediction problem in which the cloud index is forecast recursively at control-relevant temporal resolution using meteorological data available at coarser temporal resolution. This formulation directly addresses the temporal resolution mismatch identified in Chapter 1.

A baseline PINN was introduced, combining data-driven learning with partial physical knowledge of spatiotemporal cloud motion, which is incorporated through a partial differential equation governing cloud advection. The forecast target was defined as the cloud index, which serves as a physically meaningful proxy for irradiance variability. To support effective physics-based regularization, a correlation-based collocation point sampling strategy was developed that promotes physically plausible and statistically representative collocation points. Finally, input normalization and the corresponding scaling of the governing equation were described to ensure numerical stability and balanced optimization.

A recursive forecasting framework was developed to address the temporal resolution mismatch. It enables multi-step high-resolution forecasting during inference, even though models are trained on temporally sparse data. To mitigate the resulting training-inference mismatch, an improved training strategy was proposed that incorporates the recursive forecasting framework into the training process.

Finally, the validation design was established, including comparison models, evaluation metrics, and controlled experimental protocols, providing a foundation for the experimental case studies presented in the following chapter.

Chapter 4

Case study

This chapter presents the case studies used to evaluate the forecast performance of the physics-informed neural network (PINN) formulations for photovoltaic (PV) power generation forecasting in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. The simulation results are reported according to the overall validation design described in Section 3-7.

The remainder of this chapter is structured as follows. Section 4-1 describes the two experimental setups considered in this work, located in France and Hawaii. Section 4-2 details the data sources and preprocessing procedures. Section 4-3 presents the model architectures and experimental configurations. Section 4-4 presents the quantitative evaluation results for the France case study, including a detailed analysis of the baseline and improved PINN formulations. Section 4-5 reports the real-world validation results for the Hawaii case study. Section 4-6 discusses the implications of the results, and Section 4-7 summarizes the main findings of the chapter.

4-1 Experimental setups

Two complementary experimental setups are considered. The first setup, referred to as the *France* case study, constitutes the main experimental setting of this thesis. It uses temporally sparse (5 minute resolution) satellite-derived meteorological data from the Solcast service [80] for both training and evaluation, and is used to comprehensively evaluate the forecast performance of the PINN formulations for control-relevant PV power forecasting under temporal data sparsity. This case study enables the evaluation of the impact of the temporal resolution mismatch, independent of the fact that the Solcast data are satellite-derived.

The second setup, referred to as the *Hawaii* case study, provides a real-world validation of the developed recursive forecasting framework for PV systems with no local measurement infrastructure. In this realistic deployment scenario, the forecasting model has access only to temporally sparse satellite-derived data. In contrast to the France case study, this setting introduces an additional data-reality mismatch, because satellite-derived data provide only an

approximate, temporally and spatially smoothed representation of the atmospheric conditions that drive the behavior of real PV systems.

In this case study, the PINN formulations are trained on Solcast data. During inference, Solcast inputs are used to generate multi-step high-resolution forecast trajectories, which are evaluated against independent high-resolution (1 [s]) ground-truth measurements obtained from local PV monitoring infrastructure provided by the National Laboratory of the Rockies (NREL). In contrast to the France case study, in which data are available only every 5 [min], this enables the full high-resolution forecast trajectories to be evaluated directly.

4-2 Data sources and preprocessing

This section describes the data sources and the corresponding preprocessing procedures applied to construct the datasets used in the case studies. Two data sources are considered, reflecting the complementary experimental setups described in Section 4-1: coarse-resolution satellite-derived meteorological data from the Solcast service, available for both case studies, as well as high-resolution ground-truth measurements obtained from local PV monitoring infrastructure provided by the NREL, available only for the Hawaii case study.

4-2-1 Solcast coarse-resolution satellite-derived data

Meteorological data with a temporal resolution of 5 [min] were obtained from the Solcast service [80], which provides globally available estimates of meteorological variables derived from satellite imagery, numerical weather models, and data-driven post-processing.

For the France case study, Solcast data were collected for $L = 10$ measurement locations in central France. This region was chosen for its mid-latitude climate and frequent occurrence of broken-cloud regimes, which are particularly suitable for evaluating forecasting approaches at the 5 [min] temporal resolution considered, as they exhibit substantial spatial and temporal variability over such intervals [86]. The data cover the month of April 2024, representing a wide range of meteorological conditions. Nine locations were arranged on a regular spatial grid, with one additional off-grid location included to break the symmetry and encourage the model to learn a continuous spatial field rather than discrete symmetric patterns. A grid spacing of approximately 2-3 km was selected based on typical cloud motion length scales over a 5 [min] interval [66].

For the Hawaii case study, Solcast data were collected for the same measurement locations as the $L = 17$ real PV monitoring sites provided by the NREL, covering the month of March 2011. Although the geographic domain differs from the France case study, the same preprocessing procedures were applied to ensure full consistency of the data across both experimental setups.

Each Solcast record corresponds to a single spatiotemporal sample and includes, among others, the Coordinated Universal Time (UTC) reference time t_0, geographic coordinates (latitude ϕ and longitude λ), air temperature T_{air}, solar zenith angle Z, wind speed w_s and direction w_d at 100 m height, global horizontal irradiance (GHI) I, and clear-sky GHI I_{clearsky}.

Geographic coordinates were converted into local Cartesian coordinates (x, y) in meters relative to a reference location $(\phi_{\text{ref}}, \lambda_{\text{ref}})$, chosen as location 1 in the France case study and

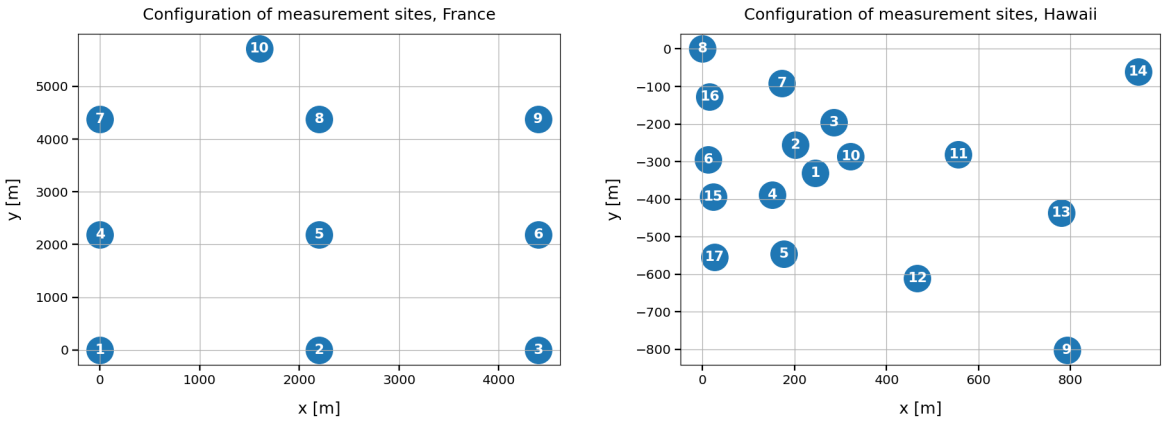

(a) France **(b)** Hawaii

Figure 4-1: Configuration of the measurement locations for the two case studies.

location 8 in the Hawaii case study. Given the limited spatial extent of the domains, a local flat-earth approximation was applied using the equirectangular projection [78]:

$$x = R\,(\gamma - \gamma_{\text{ref}})\cos(\phi_{\text{ref}}), \qquad y = R\,(\phi - \phi_{\text{ref}}), \tag{4-1}$$

where $R = 6371$ km is the mean Earth radius. The resulting spatial configurations of the measurement locations for the France and Hawaii case studies are shown in Figure 4-1, with the corresponding geographic coordinates listed in Table 4-1.

Wind speed and direction were decomposed into horizontal velocity components u and v. Following the formulation described in Section 3-1-1, the cloud index was computed as

$$c = \frac{I}{I_{\text{clearsky}}}. \tag{4-2}$$

Nighttime samples with $Z \geq 90^\circ$ were removed, as both measured and clear-sky irradiance vanish under these conditions. Samples corresponding to sunrise and sunset transitions, for which the cloud index becomes undefined, were also excluded.

The datasets for both case studies were restructured such that each sample contains the L simultaneous cloud index observations $\{c_l\}_{l=1}^{L}$. Labeled samples were then formed by defining the forecast target c_* as the cloud index at the next 5 [min] timestamp $t_* = t_0 + 300$ [s]. The forecast horizon t is explicitly included as an input variable to enable flexible horizon selection during recursive forecasting, as described in Section 3-4. The final selection of input-output variables is described in Section 3-1-1.

The resulting labeled datasets were subsequently reorganized into groups of L samples corresponding to the measurement locations, following the recursive forecasting framework developed in Section 3-4, and then randomly split into training and testing sets. Collocation points were generated following the procedure described in Section 3-2, and all variables were normalized as described in Section 3-3. This ensures that the experiments reported in this chapter are fully consistent with the methodological framework established in Chapter 3.

Table 4-1: Geographic coordinates of the measurement locations for the two case studies.

(a) France

ID	Latitude	Longitude
1	47.3486464°	0.3188°
2	47.3486464°	0.3479796°
3	47.3486464°	0.3771592°
4	47.3683464°	0.3188°
5	47.3683464°	0.3479796°
6	47.3683464°	0.3771592°
7	47.3880464°	0.3188°
8	47.3880464°	0.3479796°
9	47.3880464°	0.3771592°
10	47.4°	0.34°

(b) Hawaii

ID	Longitude	Latitude
1	21.31236°	−158.08463°
2	21.31303°	−158.08505°
3	21.31357°	−158.08424°
4	21.31183°	−158.08554°
5	21.31042°	−158.08530°
6	21.31268°	−158.08688°
7	21.31451°	−158.08534°
8	21.31533°	−158.08700°
9	21.30812°	−158.07935°
10	21.31276°	−158.08389°
11	21.31281°	−158.08163°
12	21.30983°	−158.08249°
13	21.31141°	−158.07947°
14	21.31478°	−158.07785°
15	21.31179°	−158.08678°
16	21.31418°	−158.08685°
17	21.31034°	−158.08675°

4-2-2 NREL high-resolution measurement data

Ground-truth irradiance data with a temporal resolution of 1 [s] were obtained from the PV monitoring infrastructure of the NREL [74], which provides GHI measurements recorded at $L = 17$ measurement locations and covers the month of March 2011.

The time encoding provided in the data was converted into UTC timestamps to ensure consistency with the Solcast data. Small negative irradiance values, attributed to sensor noise and calibration artifacts, were set to zero. The 1 [s] measurements were resampled to a fixed temporal resolution of t_r seconds, corresponding to the recursive forecast horizon used to generate multi-step high-resolution forecasts.

For the Hawaii case study, each sample in the Solcast test set was temporally aligned with the corresponding NREL irradiance sequence starting at the same reference time and spanning the subsequent 5 [min] interval. Each NREL irradiance trajectory was then converted element-wise into a cloud index trajectory according to Equation 3-1, where the clear-sky irradiance I_{clearsky} is obtained from the corresponding Solcast input sample and assumed constant over the 5 [min] interval. This transformation enables direct comparison with the model output and ensures consistency with the other experiments presented in this thesis.

4-3 Model architecture and experimental configuration

This section summarizes the experimental parameters used for both case studies, model architectures, training configuration and inference settings. Unless stated otherwise, the pa-

Table 4-2: Global experimental parameters used in all experiments.

Symbol	Description	Value
M	Number of collocation points	10 000
k	Number of variable clusters for collocation point sampling	3
h	KDE bandwidth for collocation point sampling	0.015
$t_{\max}$	Maximum forecast horizon	300 [s]
δ_t	Temporal margin for collocation point sampling	100 [s]
δ_s	Spatial margin for collocation point sampling	200 [m]
t_{train}	Temporal resolution of available data	300 [s]
t_r	Recursive forecast horizon	10 [s]
N_r	Number of recursive steps ($N_r = t_{\text{train}}/t_r$)	30

rameters given in Table 4-2 are used consistently across all experiments. All experiments are repeated over 10 independent random seeds, with random number generators controlled to ensure reproducibility.

Both PINN formulations and the neural networks (NNs) counterparts used in this work share a common fully-connected feedforward architecture consisting of four hidden layers with 128 neurons per layer. Model weights are initialized using the default PyTorch initialization scheme. All models are trained using the Adam optimizer with the commonly used decaying cosine annealing learning rate scheduler. During inference, all cloud index predictions are clipped to the physically meaningful range $c_* \in [0, 1]$.

To prevent the models from collapsing toward nearly uniform cloud index predictions across spatial locations and to promote the learning of meaningful spatial variability, additional mechanisms are introduced in both PINN formulations. The improved PINN is trained to predict incremental changes instead of the absolute cloud index. The baseline PINN multiplies the spatial input coordinates (x, y) by a spatial scaling factor of 2, thereby increasing the relative importance of spatial information in the training. Both mechanisms were selected based on empirical performance and were found to be optimal for the respective PINN formulations.

All experiments are evaluated on 1 000 test groups, corresponding to 10 000 individual test samples in the France case study and 17 000 samples in the Hawaii case study, reflecting the different numbers of measurement locations. The France case study uses $N = 36\,220$ training samples, while the Hawaii case study uses 53 924 training samples, in both cases obtained after excluding the test sets from the available data. In addition, $M = 10\,000$ collocation points are used. This value is selected empirically and reflects a trade-off between input domain coverage and computational cost.

The baseline PINN employs the proposed correlation-based collocation point sampling strategy. In contrast, this strategy does not yield consistent performance gains for the improved PINN, and therefore the simpler uniform sampling approach is therefore adopted for that formulation. Unless stated otherwise, the remaining model-specific hyperparameters for both PINN formulations in both case studies are listed in Table 4-3 and were selected through experimental tuning.

Table 4-3: Model-specific hyperparameters for the baseline and improved PINN formulations in the France and Hawaii case studies.

	France		Hawaii	
Description	**Baseline**	**Improved**	**Baseline**	**Improved**
Activation function	Tanh	ReLU	Tanh	ReLU
Physics-loss weight λ_{phys}	0.01	10	1	1
Initial learning rate	7×10^{-4}	1×10^{-4}	7×10^{-4}	4×10^{-4}
Final learning rate	5×10^{-5}	1×10^{-5}	7×10^{-5}	1×10^{-5}
Number of training epochs	2000	1500	2000	1500

Table 4-4: Comparison of forecast performance for all models in the France case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

Model	**RMSE**	**MAE**	**Bias**	**VE**	**Computation time [s]**
Baseline PINN	0.16436 ±0.03571	0.12008 ±0.02699	0.04244 ±0.02673	0.00128 ±0.00007	0.12 ±0.01
Baseline NN	0.43106 ±0.11806	0.32773 ±0.13084	0.32726 ±0.13084	0.00130 ±0.00006	0.26 ±0.00
Improved PINN	0.07121 ±0.00235	0.04323 ±0.00115	0.00229 ±0.00223	0.00099 ±0.00008	0.15 ±0.03
Improved NN	0.07146 ±0.00208	0.04359 ±0.00093	0.00242 ±0.00222	0.00100 ±0.00008	0.27 ±0.05
Physics-only	0.07542 ±0.00265	0.04335 ±0.00147	0.00222 ±0.00131	0.00138 ±0.00010	25.27 ±1.01
Persistence	0.07333 ±0.00249	0.04225 ±0.00130	0.00209 ±0.00155	0.00112 ±0.00009	0.16 ±0.01

4-4 Simulation results for the France case study

This section presents the evaluation results of the baseline and improved PINN formulations for control-relevant PV power generation forecasting under temporal data sparsity in the France case study. Table 4-4 summarizes the quantitative forecast performance of the PINN formulations and all comparison models, averaged over the 10 independent training seeds. All error metrics are reported for the recursively generated terminal cloud index forecasts over a 5 [min] forecast horizon and are computed with respect to the satellite-derived Solcast reference data. The errors are visualized in Figure 4-2, where the error bars indicate the standard deviation across the 10 training seeds. Because the proposed correlation-based collocation point sampling strategy did not yield consistent performance gains for the improved PINN, the simpler uniform sampling approach is adopted for that formulation, unless stated otherwise.

These results reveal several important trends. First, the baseline PINN consistently outperforms its purely data-driven NN counterpart with a reduction in RMSE of approximately 61.87%. This improvement can be attributed to the inclusion of physics-based regularization, which promotes physically plausible temporal generalization during recursive inference and

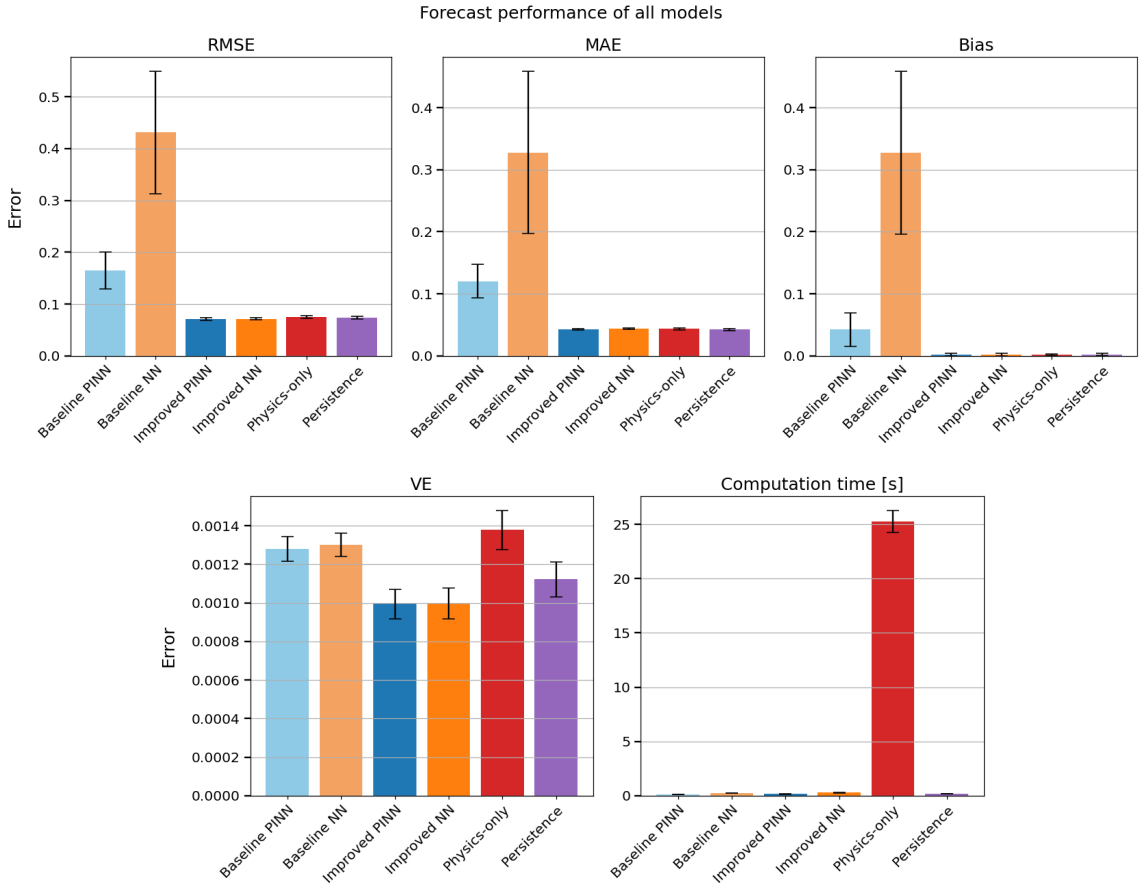

Figure 4-2: Comparison of forecast performance for all models in the France case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

mitigates error accumulation over the forecast horizon. The improved PINN further enhances the performance relative to the baseline formulation by employing the improved training strategy, resulting in a reduction in RMSE of approximately 56.67%. This strategy aligns the training objective with the recursive inference procedure by optimizing recursively generated terminal forecasts rather than single-pass forecasts. Notably, the performance gap between the improved PINN and the improved NN is minimal. This indicates that the improved training strategy benefits both models and thereby reduces the reliance on physics-based generalization to mitigate error accumulation, which was more pronounced in the baseline setting.

At the same time, the physics-only and persistence models remain highly competitive. Since cloud motion at temporal scales on the order of minutes is slowly evolving and largely advection-driven, surpassing these models over the 5 [min] forecast horizon considered in this thesis is inherently challenging. Nevertheless, the improved PINN achieves the lowest RMSE, with an average reduction of approximately 2.89% relative to persistence, which constitutes the strongest competing method. In addition, the improved PINN yields the lowest VE, indicating the most accurate preservation of spatial structure. However, the persistence model achieves the lowest MAE and bias. Finally, the physics-only model exhibits substantially higher computational cost than all other models because of the repeated evaluation of the cloud motion equation. These evaluations must be performed online during forecasting, whereas the PINN and NNs models incur their computational cost primarily during offline training, resulting in fast inference and making them more practically applicable for forecasting.

The following sections provide a detailed analysis of the forecast behavior and sensitivity of both PINN formulations under a range of experimental conditions. This analysis highlights not only the impact of physics-based regularization, but also the effectiveness of the proposed improved training strategy. Unless stated otherwise, all qualitative figures in these sections are generated from a single random seed.

4-4-1 Baseline PINN: forecast behavior and sensitivity analysis

Figure 4-3 illustrates the multi-step high-resolution forecast trajectories underlying the terminal 5 [min] forecasts for the baseline PINN and the comparison models over a temporal segment starting at 14:35 on April 16 at measurement location 1. Five consecutive 5 [min] forecast horizons are shown. The trajectories reveal distinct dynamical behaviors: the persistence forecasts remain constant by construction, the physics-only forecasts evolve in an approximately linear manner due to their reliance on fixed wind components, while both the PINN and the NN produce nonlinear, dynamically varying trajectories.

The NN exhibits pronounced deviation from the reference targets as a result of error accumulation, whereas the PINN produces forecasts that remain closer to the reference due to the inclusion of physics-based regularization. Nevertheless, consistent with the quantitative results reported in Table 4-4 and shown in Figure 4-2, the PINN does not consistently outperform the physics-only and persistence models.

The ability of the baseline PINN to capture spatial structure is illustrated in Figure 4-4, where the forecast trajectories are shown simultaneously for all measurement locations. The coordinated evolution of these trajectories reflects the updating of a shared cloud-index field

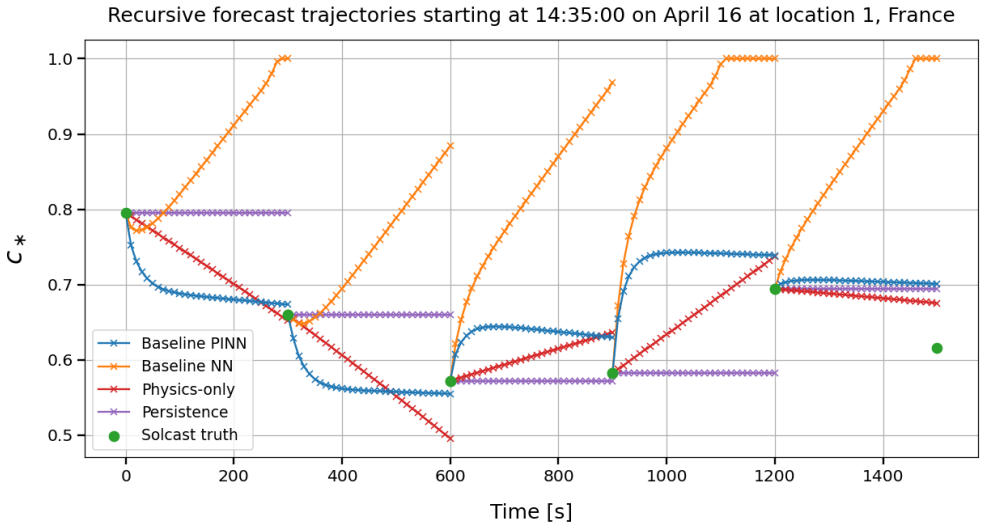

Figure 4-3: Comparison of multi-step high-resolution forecast trajectories for the baseline PINN and comparison models on April 16 at measurement location 1 in the France case study, shown over a temporal segment starting at 14:35, together with the corresponding Solcast reference targets. Five consecutive 5 [min] forecast horizons are displayed.

during recursive forecasting. However, the forecasts progressively contract toward a narrower range of spatial variability over the forecast horizon. This behavior indicates a slight conservative bias learned during training, in which the model favors reduced variability in order to minimize prediction error. Under recursive deployment, this effect compounds over successive steps, leading to a gradual loss of spatial variability until a lower bound is reached.

To isolate the contribution of key training and modeling choices to the forecast performance of the baseline PINN, a set of sensitivity experiments is conducted. These experiments evaluate the influence of the physics-loss weighting and the collocation point sampling strategy. In addition, a regime-dependent performance analysis is performed to assess the PINN's behavior under different levels of cloud variability.

Sensitivity to physics-loss weighting

To evaluate the influence of the physics-loss weight λ_{phys} on forecast performance, the baseline PINN is trained using values in the range $[10^{-3}, 10^{3}]$, while all other hyperparameters and training procedures are kept fixed. Figure 4-5 summarizes the principal error metrics for the terminal 5 [min] forecasts, averaged over the 10 independent seeds. The results indicate that moderate values of λ_{phys} provide the best balance between data fidelity and physics-based regularization, while overly weak and overly strong regularization degrade performance. For weak regularization, the PINN approaches a purely data-driven NN, which is more susceptible to error accumulation. In contrast, overly strong regularization enforces the simplified cloud motion equation too strictly, limiting the model's ability to adapt to patterns observed in the data.

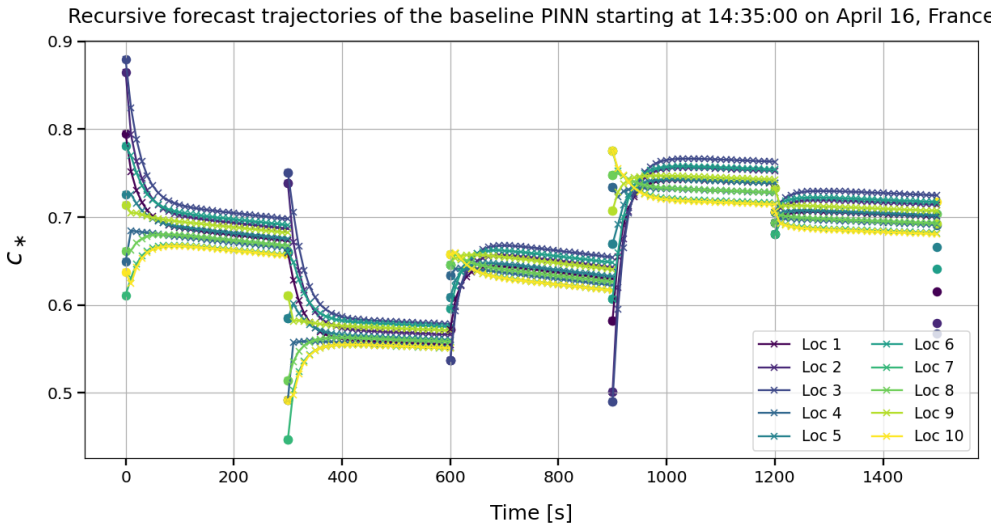

Figure 4-4: Comparison of multi-step high-resolution forecast trajectories for the baseline PINN on April 16 across all measurement locations in the France case study, shown over a temporal segment starting at $14{:}35$, together with the corresponding Solcast reference targets (indicated by colored dots for each location). Five consecutive 5 [min] forecast horizons are displayed.

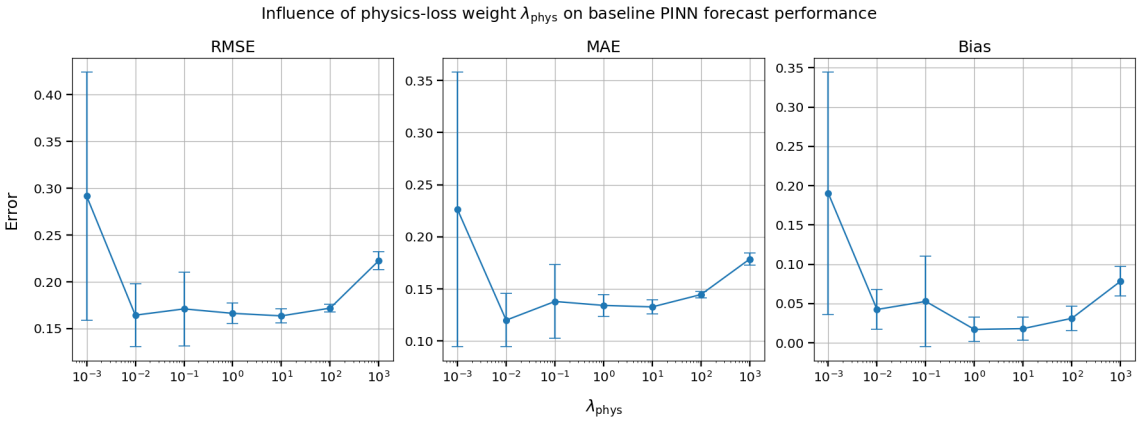

Figure 4-5: Comparison of forecast performance for the baseline PINN for different values of the physics-loss weight λ_{phys} in the France case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

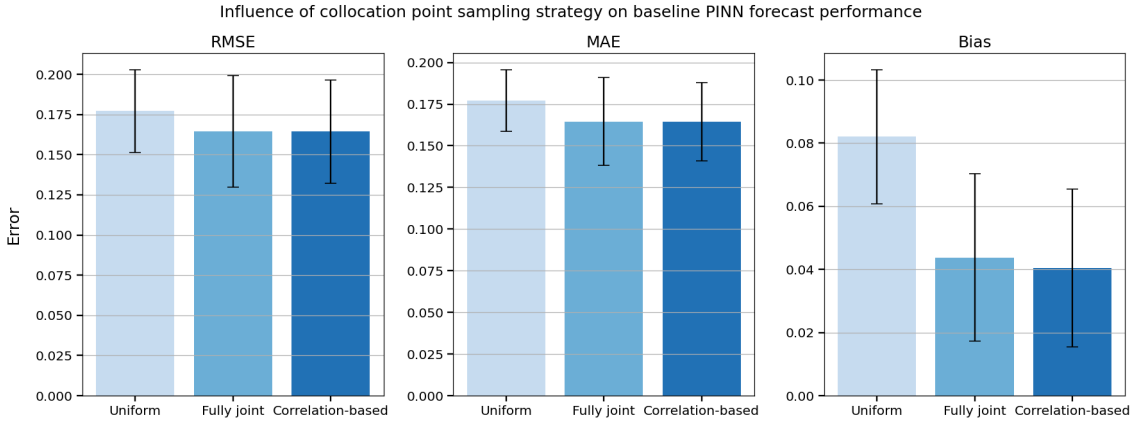

Figure 4-6: Comparison of forecast performance for the baseline PINN in the France case study using three collocation point sampling strategies: uniform sampling, fully joint KDE-based sampling, and the developed correlation-based strategy (Section 3-2). Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

Sensitivity of collocation point sampling strategy

To evaluate the influence of the collocation point sampling strategy on forecast performance, the baseline PINN is trained using three approaches: uniform sampling, fully joint KDE-based sampling, and the developed correlation-based strategy. All hyperparameters and training procedures are kept fixed. Figure 4-6 summarizes the principal error metrics for the terminal 5 [min] forecasts, averaged over the 10 independent training seeds. In this experiment, 2 000 collocation points are used, as the differences between sampling strategies become particularly pronounced in this low- to moderate-collocation regime.

The results indicate that the developed correlation-based strategy substantially outperforms uniform sampling across all metrics, with a reduction in RMSE of approximately 7.23%, demonstrating the benefit of enforcing the governing equation at combinations of input variables that are more physically plausible. Compared to fully joint KDE-based sampling, the correlation-based strategy yields only a minimal improvement. This can be explained by the resulting clustering configuration reported in Table 4-5: the largest cluster still contains $L = 10$ variables, limiting the effective dimensionality reduction of the joint sampling problem.

The resulting clusters correspond to a separation between atmospheric variables, solar geometry, and the set of cloud index observations. This clustering configuration is identical across all training seeds. The dendrogram shown in Fig. 4-7 exemplifies the underlying statistical dependency structure of the input variables.

The collocation point sampling strategy experiment is not repeated for the improved PINN, since the reduced reliance on physics-based regularization in this formulation leaves limited room for the developed correlation-based strategy to provide consistent performance gains, as indicated by preliminary experiments.

Table 4-5: Clustering configuration of the input variables used by the developed correlation-based collocation point sampling strategy.

Cluster	Variables
1	$\{c_\ell\}_{\ell=1}^{L}$
2	T_{air}, u, v
3	Z

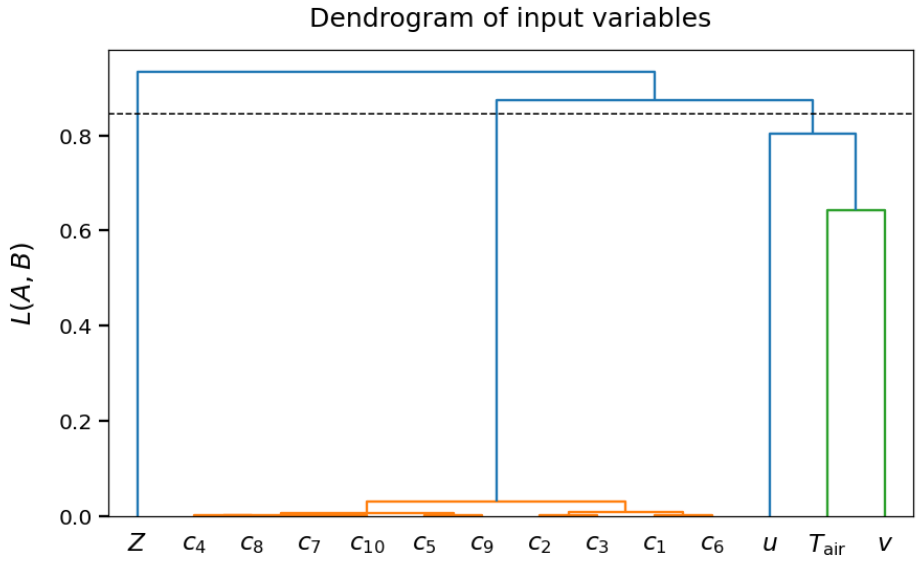

Figure 4-7: Hierarchical clustering dendrogram of the input variables used in the developed correlation-based collocation point sampling strategy. Dashed line represents the cut height.

Table 4-6: Comparison of forecast performance for the baseline PINN and persistence models across different cloud variability regimes, encoded by four quantile bins, in the France case study. RMSE values are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

c quantile bin	RMSE PINN	RMSE persistence	Win rate PINN [%]
$0-25\%$	0.08040 ± 0.05230	0.00009 ± 0.00029	0.08 ± 0.25
$25-50\%$	0.16606 ± 0.04190	0.01835 ± 0.00128	6.39 ± 2.14
$50-75\%$	0.17490 ± 0.04012	0.04652 ± 0.00176	14.96 ± 4.86
$75-100\%$	0.20661 ± 0.03819	0.13793 ± 0.00483	38.52 ± 6.98

Forecast performance under cloud variability regimes

To evaluate the forecast performance of the baseline PINN under challenging atmospheric conditions, a regime-dependent analysis is conducted across different levels of cloud variability. The test samples are grouped into quantile bins according to the magnitude of the change in cloud index over the 5 [min] forecast horizon. Within each regime, the PINN is compared against the persistence model, its strongest competitor. Table 4-6 reports the RMSE of the terminal 5 [min] forecasts for both models, averaged over the 10 independent seeds, while Figure 4-8 summarizes the corresponding PINN win rates.

As shown in Figure 4-8, the baseline PINN becomes more competitive relative to the persistence model as cloud variability intensifies. However, the win rate does not exceed the 50% threshold in any regime. Persistence therefore remains the strongest model over the 5 [min] forecast horizon. In addition, as reported in Table 4-6, the RMSE of the persistence model is consistently lower than that of the PINN across all regimes.

4-4-2 Improved PINN: forecast behavior and sensitivity analysis

Figure 4-9 shows the multi-step high-resolution forecast trajectories for the improved PINN and the comparison models on the same day (April 16) at measurement location 1. Compared to the baseline PINN, the improved PINN produces trajectories that are smoother and more closely aligned with the Solcast reference targets. The NN correspondingly shows substantially reduced error accumulation, such that only minimal differences between the two improved models remain, consistent with the quantitative results reported in Table 4-4 and Figure 4-2.

Figure 4-10 illustrates the spatial forecast behavior of the improved PINN. Compared to the baseline formulation, the improved PINN better preserves spatial variability over the forecast horizon, mitigating the progressive contraction observed in the baseline case.

Sensitivity to physics-loss weighting

The sensitivity of the improved PINN to the physics-loss weight λ_{phys} is summarized in Fig. 4-11. Compared to the baseline formulation, the improved PINN exhibits a broader range of λ_{phys} values that produce near-optimal forecast performance, indicating reduced sensitivity

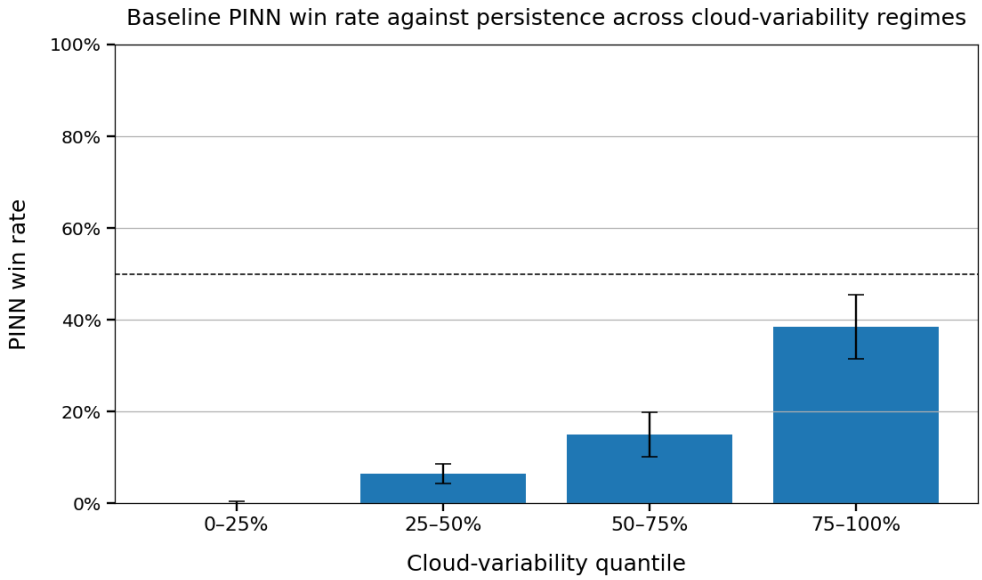

Figure 4-8: Win rate of the baseline PINN over the persistence model across different cloud variability regimes in the France case study. The win rate is based on the RMSE of the recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and is averaged over the 10 independent training seeds.

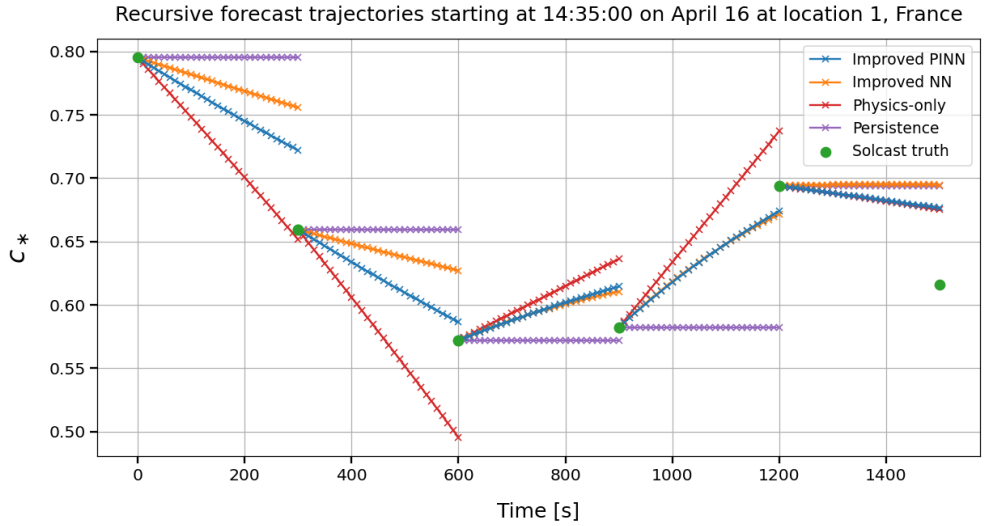

Figure 4-9: Comparison of multi-step high-resolution forecast trajectories for the improved PINN and comparison models on April 16 at measurement location 1 in the France case study, shown over a temporal segment starting at $14{:}35$, together with the corresponding Solcast reference targets. Five consecutive 5 [min] forecast horizons are displayed.

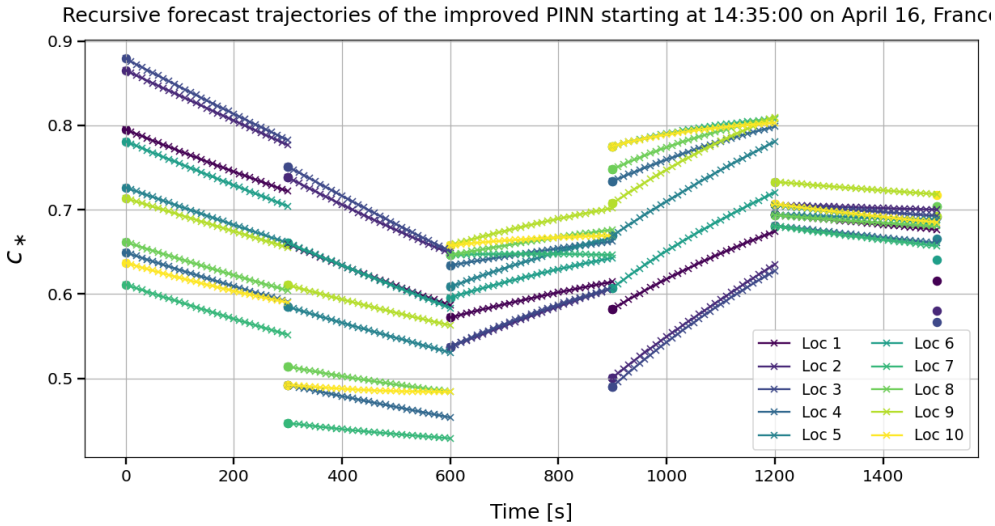

Figure 4-10: Comparison of multi-step high-resolution forecast trajectories for the improved PINN on April 16 across all measurement locations in the France case study, shown over a temporal segment starting at 14:35, together with the corresponding Solcast reference targets (indicated by colored dots for each location). Five consecutive 5 [min] forecast horizons are displayed.

to the precise choice of the physics-loss weighting. This behavior can be attributed to the improved training strategy, which reduces the reliance on physics-based regularization for mitigating error accumulation in multi-step high-resolution forecasts.

Sensitivity to training data size

To evaluate the impact of training data availability on forecast performance, the improved PINN and NN are trained using a reduced dataset with $N = 2000$ samples, while all other hyperparameters and training procedures are kept fixed. In this experiment, the correlation-based collocation point sampling strategy is employed, as it provides performance gains over uniform sampling in data-limited regimes. Figure 4-12 summarizes the principal error metrics for the terminal 5 [min] forecasts. Due to the strong sensitivity of both models to hyperparameter selection in this data-limited regime, the comparison is shown for a representative random seed after individual tuning of the physics-loss weight $\lambda_{\text{phys}} = 0.01$.

The performance of the improved NN deteriorates more severely than that of the improved PINN, with the improved PINN achieving a relative reduction in RMSE of approximately 7.3%. This indicates that physics-based regularization promotes physically plausible generalization to unseen conditions and provides a clear advantage for the PINN when training data are limited.

Forecast performance under cloud variability regimes

Figure 4-13 report the regime-dependent performance of the improved PINN. It consistently achieves higher win rates than the baseline PINN in all variability regimes and surpasses the 50% threshold in the two higher-variability bins. This demonstrates a clear advantage

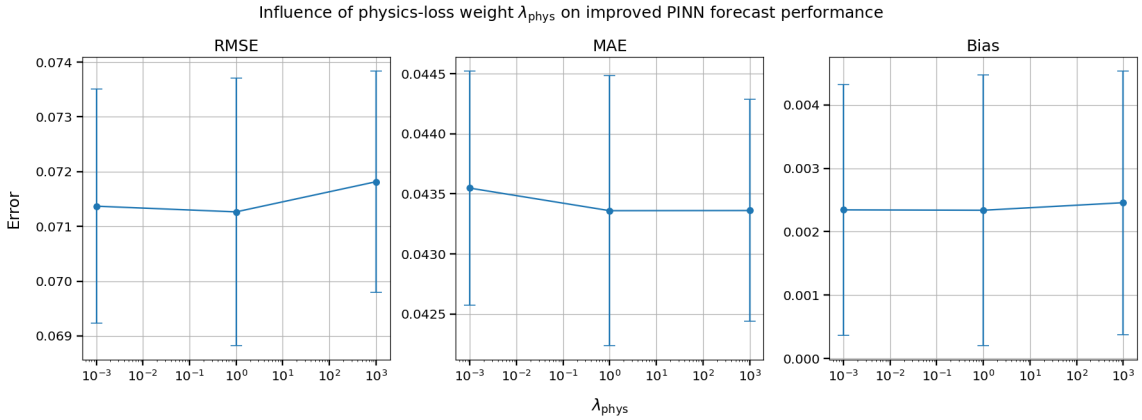

Figure 4-11: Comparison of forecast performance for the improved PINN for different values of the physics-loss weight λ_{phys} in the France case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

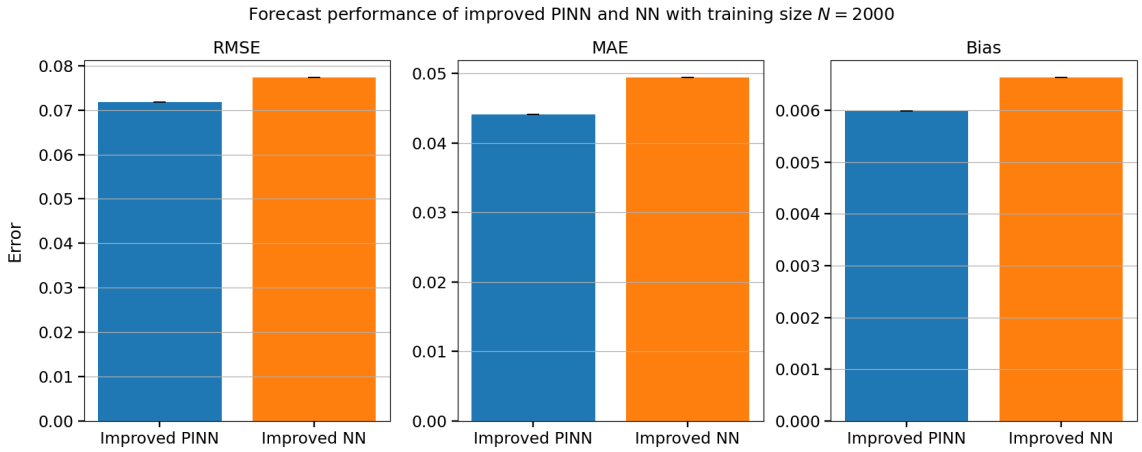

Figure 4-12: Comparison of forecast performance for the improved PINN and NN for different values of the training data size N in the France case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

Table 4-7: Comparison of forecast performance for the improved PINN and persistence models across different cloud variability regimes in the France case study. RMSE values are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and are averaged over the 10 independent training seeds.

c **quantile bin**	**RMSE PINN**	**RMSE persistence**	**Win rate PINN**[%]
$0-25\%$	0.00911 ± 0.00088	0.00009 ± 0.00029	0.00 ± 0.00
$25-50\%$	0.02582 ± 0.00134	0.01835 ± 0.00128	31.11 ± 2.94
$50-75\%$	0.04779 ± 0.00174	0.04652 ± 0.00176	53.72 ± 3.19
$75-100\%$	0.13146 ± 0.00478	0.13793 ± 0.00483	63.72 ± 3.53

over persistence under rapidly evolving atmospheric conditions. Moreover, as reported in Table 4-7, the improved PINN also achieves lower RMSE than the persistence model in the highest-variability regime, providing a relative reduction of approximately 4.69%.

4-5 Simulation results for the Hawaii case study

Unlike the France case study, the Hawaii experiment is not intended as a competitive evaluation of forecast performance among different models. Instead, it serves as a real-world validation of the developed recursive forecasting framework under a realistic PV deployment scenario. Forecast trajectories are evaluated against independent ground-truth NREL measurements, such that forecast performance is primarily constrained by the additional data-reality mismatch. The performance of the PINN formulations should therefore be interpreted as indicative of the achievable performance under these practical deployment scenarios.

Table 4-8 reports the quantitative forecast performance of the PINN formulations for the Hawaii case study, averaged over the 10 independent training seeds. The table includes the error metrics for the recursively generated terminal cloud index forecasts over a 5 [min] forecast horizon with respect to the NREL measurement data, as well as the accumulated errors over the corresponding multi-step high-resolution forecast trajectories.

The results exhibit a degradation in apparent performance relative to the France case study. While the France case study evaluates forecasts generated from satellite-derived inputs against satellite-derived reference data, the Hawaii experiment evaluates forecasts against independent high-resolution ground-truth measurements. The observed performance gap therefore primarily reflects the impact of the additional data-reality mismatch inherent to this realistic deployment scenario for real PV systems.

Figure 4-14 illustrates this effect for a temporal segment starting at 12:30 on March 27 at measurement location 1. The multi-step high-resolution forecast trajectories of the baseline and improved PINN are shown together with the corresponding high-resolution NREL measurements. The satellite-derived Solcast reference targets are also included; they provide an approximate and temporally smoothed representation of the underlying dynamics and thereby expose the additional data-reality mismatch inherent in this deployment scenario.

Although both PINN formulations remain close to the satellite-derived reference targets, the forecast trajectories neither originate from nor converge to the NREL measurements, and

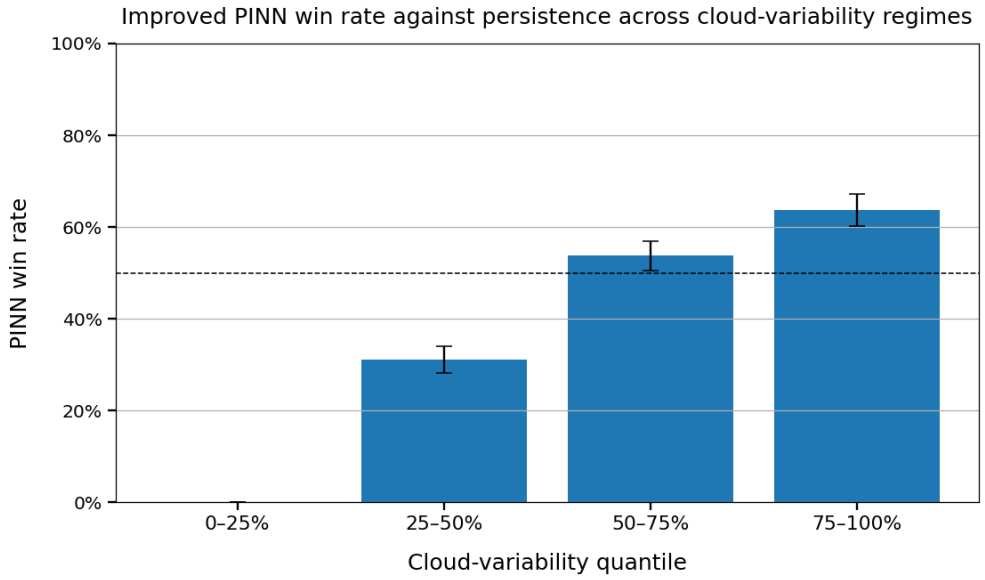

Figure 4-13: Win rate of the improved PINN over the persistence model across different cloud variability regimes, encoded by four quantile bins, in the France case study. The win rate is based on the RMSE of the recursively generated terminal 5 [min] cloud index forecasts with respect to the satellite-derived Solcast reference data and is averaged over the 10 independent training seeds.

Table 4-8: Comparison of forecast performance for the baseline and improved PINNs in the Hawaii case study. Metrics are computed on recursively generated terminal 5 [min] cloud index forecasts with respect to the NREL measurement data, as well as on the corresponding forecast trajectories. Metrics are averaged over the 10 independent training seeds.

Forecasting model	RMSE	MAE	Bias	VE
Final forecast error				
Baseline PINN	0.27326 ±0.00507	0.21558 ±0.00752	0.03795 ±0.01513	0.01496 ±0.00050
Improved PINN	0.26925 ±0.00466	0.19262 ±0.00390	0.08840 ±0.00532	0.01490 ±0.00049
Trajectory error				
Baseline PINN	0.26044 ±0.00440	0.19871 ±0.00609	0.05128 ±0.00769	0.01497 ±0.00068
Improved PINN	0.26163 ±0.00489	0.18588 ±0.00438	0.08152 ±0.00597	0.01492 ±0.00068

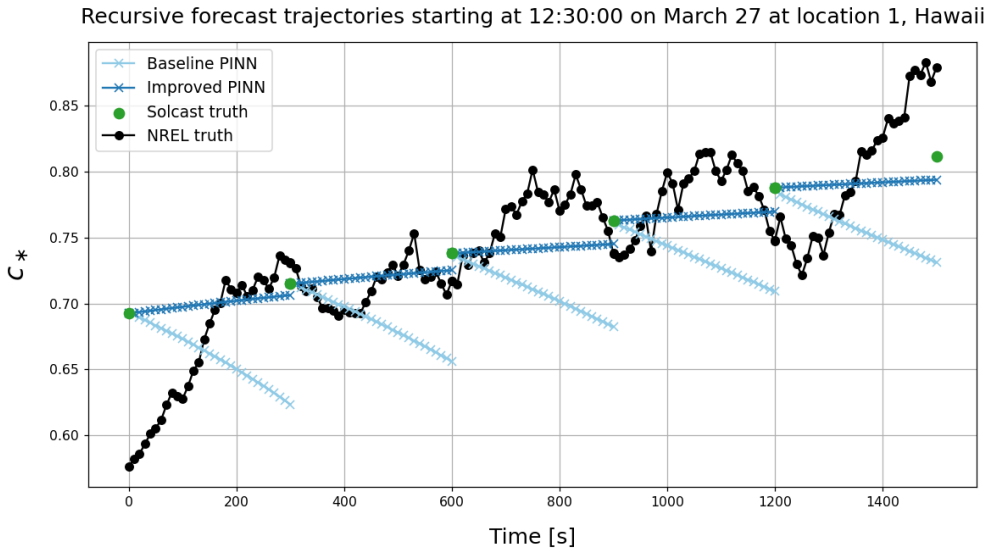

Figure 4-14: Comparison of multi-step high-resolution forecast trajectories for the baseline PINN and the improved PINN on March 27 at measurement location 1 in the Hawaii case study, shown over a temporal segment starting at 12:30, together with the corresponding NREL measurements and satellite-derived Solcast reference targets. Five consecutive 5 [min] forecast horizons are displayed.

they remain limited in reproducing the high-resolution variability observed in the measured trajectories. This indicates that the achievable forecast performance of both PINN formulations is inherently constrained by the limited fidelity of the available satellite-derived and temporally sparse data.

4-6 Discussion

The comparison between the baseline PINN and the baseline NN shows that physics-based regularization mitigates error accumulation during recursive inference. The improved training strategy reduces the reliance on physics-based regularization for this purpose by explicitly aligning the training objective with the recursive inference procedure. This explains the markedly smaller performance gap observed between the improved PINN and the improved NN. Nevertheless, the sensitivity of the improved NN to training data availability indicates that physics-based regularization remains valuable under data-limited conditions. In such regimes, it promotes physically plausible generalization to unseen conditions, resulting in the improved PINN retaining a clear data-efficiency advantage over the improved NN.

Although the physics-only and persistence models remain highly competitive over the relatively short 5 [min] horizon considered, the improved PINN offers several advantages that are relevant for power system control applications. In particular, the physics-only model becomes impractical in the present setting due to its substantially higher computational cost, arising from the repeated numerical evaluation of the cloud motion equation required to propagate the cloud field under sparse spatial observations. In contrast, the PINN provides a computationally efficient and physically plausible surrogate, making it well suited for real-time

control.

Furthermore, the improved PINN consistently outperforms the persistence model in the most challenging atmospheric conditions. These regimes are characterized by rapid and large power fluctuations in which accurate forecasting becomes particularly valuable for power system control. In addition, the relative advantage of the proposed PINN framework over persistence is expected to increase for forecast horizons longer than 5 [min]. There, cloud dynamics progressively deviate from slow-evolution assumptions and the limitations of persistence become more pronounced. This suggests that PINN-based forecasting under temporal data sparsity is well suited for power system control applications such as congestion management, which typically operate on time scales of several minutes to tens of minutes.

The improved PINN also achieves the lowest spatial VE among all models. This indicates superior preservation of spatial structure, which is important for coordinated control in distributed PV systems. In addition, this suggests that the PINN generalizes more effectively across spatial locations, enabling the transfer of learned dynamics to new spatial contexts, provided that the spatial configuration of measurement locations is similar.

The developed correlation-based collocation point sampling strategy demonstrates clear benefits over uniform sampling, while yielding only limited improvement over fully joint KDE-based sampling in the present setting. Nevertheless, the strategy is expected to offer greater benefits relative to fully joint sampling in higher-dimensional problems or in data-limited regimes.

Finally, the reduced forecast performance observed in the Hawaii case study relative to the France case study reflects the fundamentally more challenging nature of forecasting for real PV systems. In the realistic deployment scenario considered, where the available data are not only temporally sparse but also satellite-derived, an additional data-reality mismatch is introduced. This mismatch inherently constrains the achievable forecast performance of any model. Nevertheless, the developed recursive forecasting framework remains capable of producing meaningful control-relevant forecasts under such constraints, demonstrating its practical applicability for real PV systems.

4-7 Summary

This chapter presented an extensive experimental evaluation of the PINN formulations for forecasting PV power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. Two complementary case studies were considered. The France case study provided a controlled experimental setting in which a comprehensive comparison of the PINN formulations with purely data-driven NNs, a physics-only model, and a persistence model was conducted under the developed recursive forecasting framework. The Hawaii case study provides a real-world validation of this framework under a realistic PV deployment scenario in which the available data are temporally sparse and satellite-derived. This case study represents PV systems in which local measurements are unavailable. The forecasts were evaluated against independent high-resolution ground-truth measurements.

In the France case study, the improved PINN achieved the lowest RMSE among all models. Although absolute improvements were only small, the improved PINN exhibited a clear ad-

vantage under challenging atmospheric conditions with respect to persistence, demonstrated superior data efficiency relative to its NN counterpart, and showed improved computational efficiency compared to the physics-only model. Sensitivity analyses further characterized the influence of the physics-loss weighting and the collocation point sampling strategy.

The Hawaii case study highlighted the impact of additional data-reality mismatch in realistic PV deployment scenarios. Nevertheless, it demonstrated that the proposed recursive forecasting framework remains capable of producing meaningful control-relevant forecasts.

Chapter 5

Conclusions and recommendations

This thesis investigated the use of physics-informed neural networks (PINNs) for forecasting photovoltaic (PV) power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. This objective is motivated by the fundamental mismatch between the temporal resolution of commonly available meteorological data and the higher temporal resolution required by power system control applications.

This chapter first summarizes the main contributions of the thesis in Section 5-1. It then presents the main conclusions with respect to the research questions formulated in Chapter 1 in Section 5-2. Finally, Section 5-3 discusses the limitations of the present work and outlines directions for future research.

5-1 Contributions

This thesis makes the following main contributions:

- Introduced PINNs for high-resolution PV power generation forecasting (e.g., 10 [s]) in spatially distributed PV systems under temporal data sparsity (e.g., 5 [min]).
- Developed a recursive forecasting framework for generating multi-step high-resolution forecasts from models trained exclusively on temporally sparse data, together with a validation approach aligned with this temporal resolution mismatch.
- Proposed an improved training strategy that explicitly accounts for recursive deployment by incorporating the recursive forecasting framework into the training objective, thereby reducing the training-inference mismatch.
- Developed a novel correlation-based collocation point sampling strategy that generates physically plausible and statistically representative collocation points, improving the effectiveness of physics-based regularization in PINNs.

In addition, a comprehensive evaluation is conducted using two case studies, including comparisons across multiple forecasting models, collocation point sampling strategies, and training data sizes, as well as real-world validation under realistic PV deployment scenarios. Together, these contributions advance ongoing research on physics-informed learning for control-relevant renewable power generation forecasting.

5-2 Conclusions

The results of this thesis provide several important insights with respect to the research questions formulated in Chapter 1. Regarding **RQ1**, the thesis demonstrates that partial physical knowledge of atmospheric processes, represented by a simplified spatiotemporal cloud advection equation, can be effectively integrated into a PINN framework for forecasting PV power generation in spatially distributed PV systems. Although the physical model constitutes only an approximation of true cloud dynamics, its incorporation as a physics-based regularizer promotes physically plausible generalization between sparse observations.

Regarding **RQ2**, the developed recursive forecasting framework enables the generation of multi-step high-resolution forecasts from temporally sparse data. The results show that PINNs are particularly well suited for this recursive deployment, as physics-based regularization promotes physically plausible forecasts and thereby mitigates error accumulation over the forecast horizon.

With respect to **RQ3**, the improved training strategy proposed in this thesis substantially enhances multi-step high-resolution forecast performance by aligning the training objective with recursive inference. Physics-based regularization has been shown to remain valuable for regimes with limited data, as it promotes physically plausible generalization to unseen conditions. Furthermore, the PINN exhibits its greatest performance gains under challenging atmospheric conditions, where accurate forecasting is particularly valuable for power system control.

In addition, the developed correlation-based collocation point sampling strategy supports effective physics-based regularization in PINNs by promoting physically plausible and statistically representative collocation points. It outperforms conventional uniform sampling, while only marginal improvements are observed over fully joint sampling in the present setting. Nevertheless, greater benefits relative to fully joint sampling are expected in higher-dimensional problems or in data-limited regimes.

Together, these results provide a clear answer to the overarching research question of this thesis. The proposed PINN-based forecasting framework can be used effectively to forecast PV power generation in spatially distributed PV systems under temporal data sparsity and real-time power system control requirements. This is achieved by incorporating suitable partial physical knowledge of cloud dynamics into the training process to promote physically plausible generalization, deploying the models within a recursive forecasting framework to obtain multi-step forecasts at control-relevant temporal resolution, and aligning model training with this recursive deployment to further improve forecast performance.

From a broader perspective, the findings confirm that the primary value of physics-informed learning lies not necessarily in nominal operating conditions, but in precisely those practical

regimes where data are temporally sparse, uncertain, or limited, which are conditions that commonly arise in real-world power system operation.

Finally, the Hawaii case study highlighted the fundamental impact of data-reality mismatch in realistic PV deployment scenarios, where available observations are not fully representative of the true system dynamics. Importantly, for control applications, accurate forecasting of temporal variations in the cloud index is often more critical than precise estimation of its absolute magnitude. This implies that appropriate model design and training strategies, which learn the system dynamics that are relevant for control, can partially mitigate the practical impact of data-reality mismatch. Although the fidelity of the available data inevitably constrains achievable forecast performance, the proposed forecasting framework nevertheless remained capable of producing meaningful control-relevant forecasts under these conditions, providing a practical solution for PV systems worldwide that lack local measurement infrastructure.

5-3 Recommendations for future work

The results of this thesis reveal several important directions for future research. First, the physical model incorporated into the PINN is intentionally simplified and accounts only for advective cloud motion with constant wind fields over the forecast horizon. Although this is reasonable for the minute-scale forecasting horizons considered in this work [22], future research could explore the incorporation of more expressive forecasting model architectures together with richer physical representations, such as diffusive and convective dynamics or stochastic cloud evolution processes [62, 70], to further improve the realism and effectiveness of the physics-based regularization.

Furthermore, future studies could investigate the full forecasting pipeline, translating cloud index forecasts into PV power generation forecasts. In this context, the integration of additional physical models of panel dynamics, such as thermal effects and electrical response characteristics, offers a promising direction for further improving forecast performance [88].

Second, the achievable forecast performance in realistic PV deployment scenarios is inherently constrained by the fidelity of the available observations. In particular, when only temporally sparse satellite-derived data are available, an unavoidable data-reality mismatch arises between model inputs and the true atmospheric conditions at the PV system. A promising direction for future research is therefore the development of forecasting frameworks that explicitly account for these limitations. This includes data-fusion approaches that integrate complementary information sources, such as sky imagers, sparse ground-based irradiance sensors, local PV measurements, and NWP outputs, while addressing challenges such as heterogeneous temporal resolution, inconsistent spatial coverage, and measurement uncertainty. In addition, confidence-aware PINN formulations that explicitly quantify forecast uncertainty, together with online adaptation mechanisms, could further improve reliability in real deployment scenarios.

Third, integrating the proposed PINN-based forecasting framework into power system control applications, such as MPC for LFC, to evaluate the resulting control performance would enable assessing its practical value for maintaining grid stability in modern power systems.

Appendix A

Paper

Physics-informed neural networks for forecasting photovoltaic power generation for real-time control of power systems

1st D. Spee
Delft Center for Systems and Control
TU Delft
Delft, the Netherlands

2nd F. Cordiano
Delft Center for Systems and Control
TU Delft
Delft, the Netherlands

3rd A. Riccardi
Delft Center for Systems and Control
TU Delft
Delft, the Netherlands

4th B.H.K. De Schutter
Delft Center for Systems and Control
TU Delft
Delft, the Netherlands

***Abstract*—The increasing penetration of renewable energy sources into modern power systems introduces significant challenges for real-time power system control due to their inherent variability and limited controllability. Advanced control strategies such as model predictive control require multi-step forecasts at temporal resolutions on the order of seconds, while commonly available meteorological data are typically temporally sparse. This work investigates the use of physics-informed neural networks (PINNs) to forecast photovoltaic (PV) power generation in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity. The proposed PINN incorporates a spatiotemporal cloud motion equation, introducing physics-based regularization and thereby supporting physically plausible generalization between temporally sparse observations. To generate high-resolution multi-step forecasts, a recursive forecasting framework is developed, together with an improved training strategy that explicitly accounts for recursive deployment by training the model on its recursively generated forecasts. In addition, a novel correlation-based collocation point sampling strategy is developed to produce physically plausible collocation points for effective physics-based regularization. The proposed approach is evaluated on a case study in a selected region of France. The proposed PINN with improved training demonstrates superior performance under challenging atmospheric conditions compared to a persistence baseline, improved data efficiency relative to purely data-driven neural networks, and improved computational efficiency at inference compared to a physics-only model.**

I. Introduction

The global transition from centralized, dispatchable fossil-fueled power generation toward decentralized and weather-dependent renewable energy sources (RES) is fundamentally reshaping modern power systems [1]. Although this transition is essential to achieve decarbonization goals, it introduces significant challenges for reliable grid operation due to the inherent variability, limited controllability, and decentralized nature of renewable generation. Maintaining frequency stability, which is critical for the reliable operation of power systems, becomes increasingly difficult as conventional dispatchable generators are displaced by non-dispatchable renewable energy resources [2], [3]. Consequently, accurate and reliable forecasting of renewable power generation has become an important requirement for modern power system operation [4].

Advanced control strategies, such as model predictive control (MPC) [5], [6], are increasingly considered for real-time applications including load frequency control (LFC) [7], due to their ability to optimize performance under constraints. These strategies rely on multi-step forecasts of future power imbalances at temporal resolutions aligned with the control layers of interest, typically on the order of seconds for LFC.

Existing forecasting approaches struggle to meet these control requirements. In particular, physics-based models [8] provide physically plausible forecasts but require detailed knowledge of complex atmospheric processes governing renewable power generation and are computationally expensive for real-time deployment [9]. On the other hand, purely data-driven machine learning models [10] can learn complex nonlinear relationships, but depend on large volumes of data at a temporal resolution matching the required forecast resolution [9], [11]. In practice, commonly available meteorological data, including globally available satellite-derived products, are typically provided at temporal resolutions of several minutes. The temporal mismatch between data availability and control requirements poses a fundamental challenge to forecasting renewable power generation with adequate accuracy.

This challenge motivates the exploration of physics-informed machine learning approaches [12] that combine partial physical knowledge with data-driven learning to improve forecasting under temporal data sparsity. Within this framework, physics-informed neural networks (PINNs) [13] offer a promising modeling approach by embedding the governing equations directly into the training process through physics-based regularization. In principle, this enables PINNs to generalize between sparse observations and produce physically plausible forecasts at temporal resolutions higher than those

available in the data.

In addition, generating high-resolution multi-step forecasts from temporally sparse data typically relies on recursive forecasting, in which a model trained for single-step forecasting is applied recursively during inference to obtain a forecast trajectory. This recursive strategy is known to suffer from error accumulation [14], [15], further motivating the investigation of physics-based regularization as a means of mitigating this effect under recursive deployment. Nevertheless, the practical effectiveness of PINNs for generating multi-step renewable power forecasts at temporal resolutions on the order of seconds under temporal data sparsity remains largely unexplored.

This work focuses on photovoltaic (PV) power generation in spatially distributed PV systems. At intra-hour time scales, PV power generation variability is dominated by spatiotemporal cloud motion, which induces rapid fluctuations in solar irradiance that propagate across spatially distributed PV systems [16]. Moreover, many PV systems lack local sensing infrastructure and must rely exclusively on temporally sparse satellite-derived data, making PV systems a representative and practically important case for investigating forecasting under temporal data sparsity.

This paper makes the following main contributions:

- It introduces PINNs for multi-step PV power generation forecasting in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity.
- It develops a recursive forecasting framework for PINNs that enables multi-step high-resolution forecasting from models trained exclusively on temporally sparse data, together with a validation approach aligned with this temporal resolution mismatch.
- It proposes an improved training strategy for PINNs that explicitly accounts for recursive deployment by training the model on its recursively generated forecasts, thereby reducing the training-inference mismatch.
- It develops a novel correlation-based collocation point sampling strategy for PINNs that generates physically plausible and statistically representative collocation points, improving the effectiveness of physics-based regularization.

The remainder of this paper is organized as follows. Section II presents the proposed PINN-based forecasting methodology, including the recursive forecasting framework, the improved training strategy, and the correlation-based collocation point sampling strategy. Section III describes the experimental setup of the case study and discusses the results. Finally, Section IV summarizes the main findings of this work and outlines directions for future work.

II. Methodology

This section presents the proposed PINN-based framework for multi-step PV power generation forecasting in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity.

TABLE I
Overview of input and output variables.

	Description	Time	Location	Unit
Inputs				
(x, y)	Target location	-	-	m
T_{air}	Air temperature	t_0	(x, y)	°C
Z	Solar zenith angle	t_0	(x, y)	°
(u, v)	Wind components	t_0	(x, y)	m/s
$\{c_l\}_{l=1}^{L}$	Cloud index observations	t_0	L sites	-
t	Forecast horizon	-	-	s
Output				
c_*	Target cloud index	t_*	(x, y)	-

A. PINN formulation

At intra-hour time scales, variability in PV power generation is dominated by spatiotemporal cloud motion [16]. Therefore, the PINN forecasts the cloud index

$$c = \frac{I}{I_{\text{clearsky}}}, \tag{1}$$

where I denotes the global horizontal irradiance (GHI) and I_{clearsky} denotes the corresponding clear-sky GHI. The cloud index isolates irradiance fluctuations caused by cloud cover and serves as a physically meaningful proxy for PV power generation [17]–[19].

The PINN output approximates the mapping

$$c_\theta : \mathbf{q} \mapsto c_*, \tag{2}$$

with input vector

$$\mathbf{q} = [x,\ y,\ t,\ T_{\text{air}},\ Z,\ u,\ v,\ \{c_l\}_{l=1}^{L}], \tag{3}$$

where (x, y) denote the target location, t is the forecast horizon, T_{air} is the air temperature, Z (°) is the solar zenith angle, (u, v) are the wind components, and $\{c_l\}_{l=1}^{L}$ are the cloud index observations at L measurement locations. All inputs are observed at reference time t_0, while the output corresponds to time $t_* = t_0 + t$. An overview of all variables, including their temporal and spatial interpretation and units, is provided in Table I.

Partial physical knowledge is incorporated into the PINN through the cloud advection equation [20], [21]

$$\frac{\partial c}{\partial t} + u\frac{\partial c}{\partial x} + v\frac{\partial c}{\partial y} = 0, \tag{4}$$

which describes the transport of the cloud index through space and time driven by the wind velocity components (u, v), which are assumed constant over the forecast horizon.

The PINN training minimizes a composite loss

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}}\,\mathcal{L}_{\text{phys}}, \tag{5}$$

where $\lambda_{\text{phys}} > 0$ is a weighting parameter that controls the strength of the physics-based regularization.

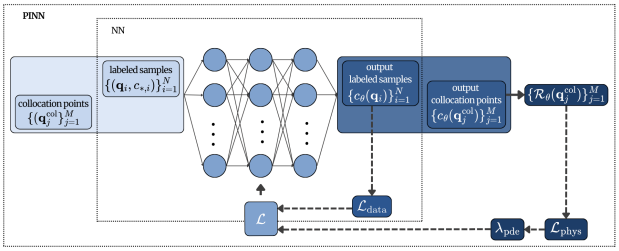

Fig. 1. Overview of the PINN training framework.

The data loss penalizes discrepancies between predictions $c_\theta(\mathbf{q})$ and observations c_*

$$\mathcal{L}_{\text{data}} = \frac{1}{N}\sum_{i=1}^{N}\left(c_\theta(\mathbf{q}_i) - c_{*,i}\right)^2, \tag{6}$$

where $\{(\mathbf{q}_i, c_{*,i})\}_{i=1}^{N}$ denotes the set of N labeled training samples.

The physics loss penalizes violations of the cloud motion equation

$$\mathcal{L}_{\text{phys}} = \frac{1}{M}\sum_{j=1}^{M}\left(\mathcal{R}_\theta(\mathbf{q}_j^{\text{col}})\right)^2, \tag{7}$$

where $\{\mathbf{q}_j^{\text{col}}\}_{j=1}^{M}$ denotes the set of M collocation points. The residual of the cloud motion equation is defined as

$$\mathcal{R}_\theta(\mathbf{q}^{\text{col}}) = \frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial t} + u\frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial x} + v\frac{\partial c_\theta(\mathbf{q}^{\text{col}})}{\partial y}. \tag{8}$$

All partial derivatives are computed using automatic differentiation [22]. An overview of the PINN training framework is illustrated in Fig 1.

B. Normalization and scaling

Proper normalization of input variables is essential for numerical stability and effective training of PINNs [23]. All input variables, except the cloud index observations $\{c_l\}_{l=1}^{L}$ which are already bounded in $[0, 1]$, are normalized using min-max scaling based on the training data range. For a generic variable q, normalization is given by

$$q' = \frac{q - q_{\min}}{\Delta q}, \tag{9}$$

where $\Delta q = q_{\max} - q_{\min}$. Both the observed training samples and the collocation points are normalized using this transformation, resulting in scaled inputs $\mathbf{q}'$ and $\mathbf{q}^{\text{col}'}$, respectively.

Accordingly, the cloud motion equation is reformulated in the normalized space using the chain rule. This yields a dimensionless equation of the form

$$\frac{\partial c}{\partial t'} + \tilde{u}\frac{\partial c}{\partial x'} + \tilde{v}\frac{\partial c}{\partial y'} = 0, \tag{10}$$

where $\tilde{u} = u\,\Delta t/\Delta x$ and $\tilde{v} = v\,\Delta t/\Delta y$. The physical wind components (u, v) are recovered from their normalized counterparts through the inverse min-max transformation

$$u = u_{\min} + u'\,\Delta u, \quad v = v_{\min} + v'\,\Delta v. \tag{11}$$

This scaled formulation of the cloud motion equation produces a well-conditioned physics loss that remains numerically compatible with the data loss during training.

C. Correlation-based collocation point sampling

Collocation points determine where the governing equation is enforced during training and therefore play a central role in the forecast performance of PINNs. In PV power generation forecasting, the input space consists of multiple meteorological variables that are both statistically dependent and physically coupled. Independent sampling of these variables can generate physically implausible combinations, while fully joint sampling is often poorly representative of the true underlying distribution in high-dimensional spaces under limited data [24]. Both effects degrade the effectiveness of physics-based regularization. This paper therefore proposes a correlation-based sampling strategy that preserves dominant statistical correlations while avoiding high-dimensional joint sampling. The approach consists of four main steps: correlation analysis, variable clustering, cluster-wise joint sampling, and spatiotemporal sampling.

Let $\mathbf{Q} \in \mathbb{R}^{n\times d}$ denote the matrix of the labeled training samples, where each column $\mathbf{q}_j$ corresponds to one of the d input variables included in the correlation analysis. Statistical dependencies between variables q_j and q_m are quantified using the Pearson correlation coefficient [25]

$$\rho_{jm} = \frac{\operatorname{cov}(q_j, q_m)}{\sigma_{q_j}\sigma_{q_m}}. \tag{12}$$

This is transformed into a dissimilarity measure

$$D_{jm} = 1 - |\rho_{jm}|, \tag{13}$$

which is used for hierarchical agglomerative clustering [26] with the average linkage criterion [27] to identify clusters of strongly correlating variables.

Clusters are sampled independently, while within each cluster with dimension d_c, collocation points are generated by joint sampling using multivariate kernel density estimation (KDE) [28]. Given the training samples $\{\mathbf{q}_i\}_{i=1}^{n} \subset \mathbb{R}^{d_c}$, the estimated density is

$$\hat{f}(\mathbf{q}) = \frac{1}{nh^{d_c}}\sum_{i=1}^{n} K\left(\frac{\mathbf{q} - \mathbf{q}_i}{h}\right), \tag{14}$$

with bandwidth parameter $h > 0$ and Gaussian kernel

$$K(\mathbf{z}) = (2\pi)^{-d_c/2}\exp\left(-\frac{1}{2}\|\mathbf{z}\|^2\right). \tag{15}$$

Samples are filtered to enforce basic physical constraints, including daylight conditions and $0 < c_l \leq 1$. Spatial coordinates are sampled uniformly over the buffered convex hull of the measurement locations, and forecast horizons are sampled uniformly over the training interval. This yields physically plausible and statistically representative collocation points, supporting effective physics-based regularization during PINN training.

D. Recursive forecasting framework

The available training data are temporally sparse, with resolution t_{train}, while real-time control requires forecasts at much finer resolution t_r. To bridge this mismatch, this paper develops a recursive forecasting framework that enables multi-step high-resolution forecasts from models trained exclusively on temporally sparse data.

At reference time t_0, the information available about the spatially distributed PV system is represented by a set of L input vectors

$$\{\mathbf{q}_\ell(t_0)\}_{\ell=1}^{L}, \tag{16}$$

where each $\mathbf{q}_\ell = [\mathbf{m}_\ell(t_0), \mathbf{c}(t_0)]$ contains local auxiliary variables $\mathbf{m}_\ell(t_0)$, which are assumed to remain constant over the forecast horizon and include spatial coordinates, the forecast horizon and meteorological variables, as well as the shared cloud index field

$$\mathbf{c}(t_0) = \{c_j(t_0)\}_{j=1}^{L}. \tag{17}$$

Importantly, the cloud index field $\mathbf{c}(t_0)$ is identical across all input vectors and therefore constitutes a shared state. Updating this shared state requires the simultaneous generation of cloud index forecasts at all L locations. The trained model c_θ is therefore applied jointly across all locations to forecast the cloud field at the first recursive step

$$\hat{\mathbf{c}}(t_0 + t_r) = \{c_\theta(\mathbf{q}_\ell(t_0))\}_{\ell=1}^{L}. \tag{18}$$

This is used to update the inputs as

$$\mathbf{q}_\ell(t_0 + t_r) = [\mathbf{m}_\ell(t_0), \hat{\mathbf{c}}(t_0 + t_r)]. \tag{19}$$

Repeating this procedure for a desired number of steps $k = 1, \ldots, N_r$ yields a recursive multi-step forecast

$$\hat{\mathbf{c}}(t_0 + kt_r) = \{c_\theta(\mathbf{q}_\ell(t_0 + (k-1)t_r))\}_{\ell=1}^{L}. \tag{20}$$

Because no new observations are available during recursion, forecast errors accumulate over successive steps. However, because PINNs incorporate physical constraints during training, they promote physically plausible generalization and can mitigate this effect under recursive deployment.

Because the available data are temporally sparse, intermediate recursive forecasts cannot be directly validated. Validation is therefore performed by executing the recursive forecasting procedure over a forecast horizon equal to the training resolution t_{train}, corresponding to $N_r = t_{\text{train}}/t_r$ recursive steps, and comparing the resulting terminal forecast $\hat{\mathbf{c}}(t_0 + N_r\, t_r)$ with the observed cloud field $\mathbf{c}(t_0 + t_{\text{train}})$. This terminal comparison evaluates the accumulated error over the forecast horizon and provides a measure of multi-step high-resolution forecast performance under temporal data sparsity.

E. Recursive training strategy

Training on temporally sparse data using single-step supervision creates a mismatch with inference-time recursive deployment. To address this, this paper proposes an improved training strategy that explicitly incorporates recursive forecasting into the training objective.

Each training sample i corresponding to a reference time t_i consists of L input vectors

$$\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L} = [\mathbf{m}_\ell(t_i), \mathbf{c}(t_i)], \tag{21}$$

and the corresponding observed cloud field target

$$\mathbf{c}_{*,i} = \mathbf{c}(t_i + t_{\text{train}}). \tag{22}$$

During training, the model is recursively applied for N_r steps, defining the rollout operator

$$\mathcal{C}_\theta^{(N_r)}\big(\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L}\big) = \hat{\mathbf{c}}(t_i + N_r t_r). \tag{23}$$

The recursive data loss is then given by

$$\mathcal{L}_{\text{data}}^{\text{rec}} = \frac{1}{NL}\sum_{i=1}^{N}\left\|\mathcal{C}_\theta^{(N_r)}\big(\{\mathbf{q}_{i,\ell}\}_{\ell=1}^{L}\big) - \mathbf{c}_{*,i}\right\|_2^2. \tag{24}$$

This training objective aligns model optimization with inference-time deployment, reducing error accumulation and improving multi-step forecast performance under temporal data sparsity.

III. Case Study results

This section evaluates the proposed PINN-based forecasting framework for multi-step PV power forecasting in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity.

A. Comparison models and validation design

Two PINN formulations are evaluated: a baseline PINN trained with single-pass supervision, and an improved PINN trained with the improved strategy based on recursive forecasting. They are compared with a persistence model, which assumes constant behavior over the forecast horizon, purely data-driven neural networks (NNs) with and without the improved training strategy, and a physics-only cloud advection model. The latter uses the closed-form solution of the cloud advection equation [29],

$$c(x, y, t + t_r) = c(x - u\, t_r,\ y - v\, t_r,\ t), \tag{25}$$

which translates the current cloud field forward in time along the wind direction.

All models are evaluated using identical datasets, network architectures, and recursive forecasting procedures. Forecast performance is quantified by comparing terminal recursive cloud index field forecasts $\hat{\mathbf{c}}$ with corresponding observed cloud index field targets $\mathbf{c}_*$. The following error metrics are used:

- **Root mean squared error (RMSE)**:

$$\text{RMSE} = \sqrt{\frac{1}{NL}\sum_{i=1}^{N}\|\hat{\mathbf{c}}_i - \mathbf{c}_{*,i}\|_2^2}. \tag{26}$$

- **Mean absolute error (MAE)**:

$$\text{MAE} = \frac{1}{NL}\sum_{i=1}^{N}\|\hat{\mathbf{c}}_i - \mathbf{c}_{*,i}\|_1. \tag{27}$$

- **Bias**:

$$\text{Bias} = \frac{1}{NL} \sum_{i=1}^{N} \sum_{l=1}^{L} \left(\hat{c}_{i,l} - c_{*,i,l} \right). \tag{28}$$

- **Variability error (VE)**:

$$\text{VE} = \frac{1}{N} \sum_{i=1}^{N} \left| \text{Var}(\hat{\mathbf{c}}_i) - \text{Var}(\mathbf{c}_{*,i}) \right|, \tag{29}$$

where $\text{Var}(\cdot)$ denotes the variance across the L spatial measurement locations within a forecast instance. These metrics are reported in terms of the mean and standard deviation over 10 independent training seeds.

B. Experimental setup and data

The proposed framework is evaluated using a case study in central France based on temporally sparse satellite-derived meteorological data obtained from the Solcast service [30] with a temporal resolution of $t_{\text{train}} = 5$ min. The data cover April 2024 and include measurements at $L = 10$ spatially distributed locations. The selected region exhibits frequent broken-cloud regimes, which induce strong spatiotemporal variability and are therefore well suited for evaluating control-relevant PV power generation forecasting [31]. The measurement locations are arranged on an approximately 2–3 km grid, consistent with typical cloud advection length scales over t_{train} [32].

After removing nighttime samples, the cloud index is computed for each remaining data record according to (1). The corresponding forecast target c_* is defined as the cloud index at the next time instant $t_* = t_0 + t_{\text{train}}$. Training samples are constructed such that each sample contains the simultaneous cloud index observations from all L locations, representing the spatial cloud field at the reference time t_0. High-resolution multi-step forecasts are then generated using the recursive forecasting framework described in Section II, for which training samples are grouped into sets of L input vectors. A control-relevant temporal resolution of $t_r = 10$ s is selected, corresponding to $N_r = t_{\text{train}}/t_r = 30$ recursive steps over the forecast horizon.

C. Model Architecture and Training Configuration

All NN-based models use the same fully-connected feed-forward architecture with four hidden layers of 128 neurons each. The baseline PINN uses Tanh activation functions, while the improved PINN uses ReLU activations, both empirically determined. Models are trained using the Adam optimizer with a cosine annealing learning-rate scheduler.

To promote the learning of meaningful spatial variability and avoid degenerate solutions with nearly uniform cloud index predictions, additional mechanisms are introduced in both PINN formulations. The improved PINN is trained to predict incremental changes in the cloud index rather than absolute values. The baseline PINN scales the spatial input coordinates (x, y) by a factor of two, increasing the relative influence of spatial information during training. These design choices were empirically selected and are fixed across all experiments.

In addition, the proposed correlation-based collocation point sampling strategy is tested for both PINNs.

Unless stated otherwise, all experiments use the hyperparameters listed in Table II.

TABLE II
TRAINING DATA, COLLOCATION POINT SAMPLING, AND MODEL HYPERPARAMETERS.

Parameter	Value
Labeled data and collocation points	
Number of training samples N	36 220
Number of test samples	10 000
Number of collocation points M	10 000
KDE bandwidth h	0.015
Number of clusters k	3
Baseline PINN	
Physics-loss weight λ_{phys}	0.01
Initial learning rate	7×10^{-4}
Final learning rate	5×10^{-5}
Training epochs	2000
Improved PINN	
Physics-loss weight λ_{phys}	10
Initial learning rate	1×10^{-4}
Final learning rate	1×10^{-5}
Training epochs	1500

D. Overall quantitative performance

Table III summarizes the quantitative forecast performance of all models on recursively generated terminal cloud-index forecasts over the 5 min forecast horizon. Because the proposed correlation-based collocation point sampling strategy did not yield consistent performance gains for the improved PINN, the simpler uniform sampling approach is adopted for that formulation, unless stated otherwise.

The baseline PINN substantially outperforms its purely data-driven NN counterpart, achieving a reduction in RMSE of approximately 61.9%. This improvement reflects the benefit of physics-based regularization, which promotes physically plausible generalization between sparse observations and mitigates error accumulation during recursive inference. The improved PINN further reduces the RMSE by approximately 56.7% relative to the baseline PINN through the proposed recursive training strategy. Notably, the performance gap between the improved PINN and the improved NN is small, indicating that the improved training strategy benefits both models by explicitly accounting for error accumulation.

The physics-only and persistence models remain highly competitive over the 5 min horizon, reflecting the slow and predominantly advection-driven cloud dynamics at this time scale. Nevertheless, the improved PINN achieves the lowest overall RMSE, providing an average improvement of 2.8% relative to persistence. In addition, it achieves the lowest spatial variability error, demonstrating superior preservation of spatial structure.

TABLE III
COMPARISON OF FORECAST PERFORMANCE FOR ALL MODELS FOR RECURSIVELY GENERATED TERMINAL 5 MIN CLOUD INDEX FORECASTS.

Model	RMSE	MAE	Bias	VE	Computation time (s)
Baseline PINN	0.16436 ± 0.03571	0.12008 ± 0.02699	0.04244 ± 0.02673	0.00128 ± 0.00007	0.12 ± 0.01
Baseline NN	0.43106 ± 0.11806	0.32773 ± 0.13084	0.32726 ± 0.13084	0.00130 ± 0.00006	0.26 ± 0.00
Improved PINN	0.07121 ± 0.00235	0.04323 ± 0.00115	0.00229 ± 0.00223	0.00099 ± 0.00008	0.15 ± 0.03
Improved NN	0.07146 ± 0.00208	0.04359 ± 0.00093	0.00242 ± 0.00222	0.00100 ± 0.00008	0.27 ± 0.05
Physics-only	0.07542 ± 0.00265	0.04335 ± 0.00147	0.00222 ± 0.00131	0.00138 ± 0.00010	25.27 ± 1.01
Persistence	0.07333 ± 0.00249	0.04225 ± 0.00130	0.00209 ± 0.00155	0.00112 ± 0.00009	0.16 ± 0.01

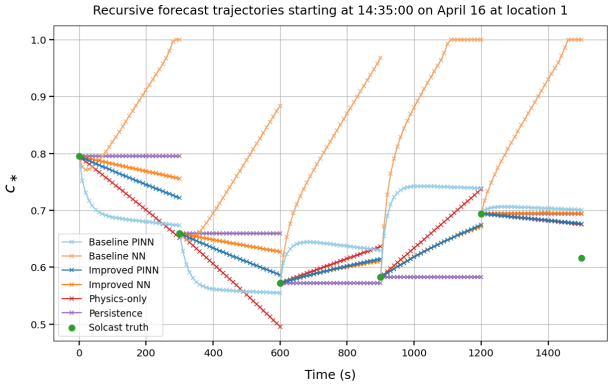

Fig. 2. Representative multi-step high-resolution forecast trajectories for all models on April 16 at measurement location 1 over five consecutive 5 min forecast horizons.

Moreover, the physics-only model incurs substantially higher computational cost due to repeated online evaluations of the cloud motion equation, whereas the learned models provide fast inference after offline training.

E. Qualitative forecast behavior

Fig. 2 illustrates representative multi-step high-resolution forecast trajectories for all models on April 16 at measurement location 1 over five consecutive 5 min forecast horizons. The trajectories reveal distinct dynamical behaviors: persistence remains constant by construction, the physics-only model evolves approximately linearly due to its reliance on fixed wind components, while both PINN and NN models produce nonlinear, dynamically varying trajectories.

The baseline NN exhibits pronounced error accumulation, while the baseline PINN mitigates this effect through physics-based regularization. The improved PINN further enhances performance by producing smoother trajectories that are more directly targeted at the reference signal.

F. Collocation point sampling analysis

To evaluate the impact of the proposed correlation-based collocation point sampling strategy, the baseline PINN is trained using three approaches: uniform sampling, fully joint KDE-based sampling, and the correlation-based strategy. The comparison is conducted in a low- to moderate-collocation regime with $M = 2\,000$ points, where the impact of sampling choices is more pronounced.

TABLE IV
EFFECT OF COLLOCATION POINT SAMPLING STRATEGY ON FORECAST PERFORMANCE FOR THE BASELINE PINN FOR RECURSIVELY GENERATED TERMINAL 5 MIN CLOUD INDEX FORECASTS.

Sampling Strategy	RMSE	MAE	Bias
Uniform	0.17719 $\pm$0.02576	0.12899 $\pm$0.01854	0.08204 $\pm$0.02129
Fully joint	0.16449 $\pm$0.03490	0.12008 $\pm$0.02645	0.04380 $\pm$0.02654
Correlation-based	0.16438 $\pm$0.03223	0.12000 $\pm$0.02359	0.04044 $\pm$0.02500

Table IV shows that the correlation-based strategy consistently outperforms uniform sampling across all reported metrics, achieving an RMSE reduction of approximately 7.2%, demonstrating the benefit of enforcing the governing equation at physically plausible collocation points. In contrast, only marginal gains are observed relative to fully joint sampling. This is explained by the clustering configuration produced by the correlation-based approach, in which the largest cluster still contains all $L = 10$ cloud index variables $\{c_l\}_{l=1}^{L}$, thereby limiting the effective reduction of the joint sampling dimensionality.

For the improved PINN, no consistent performance gains from the proposed sampling strategy are observed, as the improved training strategy substantially reduces the PINN's reliance on physics-based regularization to mitigate error accumulation.

G. Performance of the improved PINN under challenging conditions

This section evaluates the forecast performance of the improved PINN under two challenging regimes: rapidly evolving atmospheric conditions, which are critical for real-time power system control, and limited training data availability, which stresses the model's ability to generalize to unseen conditions.

1) Challenging weather conditions: To evaluate performance under rapidly evolving atmospheric conditions, a regime-dependent analysis based on cloud variability is performed. Test samples are grouped into quantile bins according to the magnitude of the change in cloud index over the 5 min forecast horizon. Within each regime, the improved PINN is compared with the persistence model, which constitutes the strongest comparison model in this case study.

The results are summarized in Table V and show that the improved PINN exceeds the 50% win-rate threshold in the

TABLE V
FORECAST PERFORMANCE OF THE IMPROVED PINN AND THE PERSISTENCE MODEL ACROSS CLOUD VARIABILITY REGIMES FOR RECURSIVELY GENERATED TERMINAL 5 MIN CLOUD INDEX FORECASTS.

Quantile bin	RMSE (PINN)	RMSE (Persistence)	Win rate (%)
0–25%	0.00911 ±0.00088	0.00009 ±0.00029	0.00 ±0.00
25–50%	0.02582 ±0.00134	0.01835 ±0.00128	31.11 ±2.94
50–75%	0.04779 ±0.00174	0.04652 ±0.00176	53.72 ±3.19
75–100%	0.13146 ±0.00478	0.13793 ±0.00483	63.72 ±3.53

TABLE VI
FORECAST PERFORMANCE OF THE IMPROVED PINN AND THE IMPROVED NN UNDER DATA-LIMITED CONDITIONS FOR RECURSIVELY GENERATED TERMINAL 5 MIN CLOUD INDEX FORECASTS.

Forecasting model	RMSE	MAE	Bias
Improved PINN	0.07172	0.04414	0.00599
Improved NN	0.07737	0.04947	0.00664

two highest-variability regimes. In the most upper regime, the improved PINN also achieves a lower RMSE than the persistence model, with a relative reduction of approximately 4.69%. These findings indicate that the proposed PINN provides a clear advantage in challenging weather conditions.

2) Data-limited conditions: To evaluate performance in data-limited regimes, the improved PINN and NN are trained using a reduced dataset of $N = 2000$ samples. In this experiment, the correlation-based collocation point sampling strategy is employed, as it provides performance gains over uniform sampling when training data are limited.

The results are reported in Table VI. Due to the sensitivity of both models to hyperparameter selection in this data-limited regime, the comparison is shown for a representative random seed after individual tuning of the physics-loss weight $\lambda_{\text{phys}} = 0.01$. The performance of the improved NN deteriorates more severely than that of the improved PINN, with the improved PINN achieving a relative reduction in RMSE of approximately 7.3%. This indicates that physics-based regularization promotes physically plausible generalization to unseen conditions and provides a clear advantage for the PINN when training data are limited.

H. Discussion

Although physics-only and persistence models remain highly competitive over the relatively short 5 min forecast horizon considered, the improved PINN offers important advantages for power system control. In particular, the physics-only model becomes impractical due to its high computational cost associated with repeated numerical propagation of the cloud field, whereas the PINN provides a computationally efficient and physically plausible surrogate suitable for real-time deployment. Moreover, the improved PINN consistently outperforms persistence in the most challenging atmospheric conditions characterized by rapid cloud evolution. These regimes are precisely those in which accurate forecasting is most valuable for real-time control. The relative advantage of the proposed framework is therefore expected to increase for forecast horizons longer than 5 min, where cloud dynamics progressively deviate from slow-evolution assumptions and persistence degrades more rapidly.

The improved PINN also achieves superior preservation of spatial variability, as reflected by the lowest variability error among all models. This property is particularly relevant for coordinated control of spatially distributed PV systems and indicates improved generalization across spatial locations, enabling transferability to new spatial configurations with similar geometry.

Finally, the proposed correlation-based collocation sampling strategy demonstrates clear benefits over uniform sampling and provides a systematic approach to generate physically plausible collocation points, supporting effective physics-based regularization in PINN training. Although only modest improvements over fully joint sampling are observed in the present setting, the approach is expected to become increasingly valuable in higher-dimensional problems and in data-limited regimes.

IV. CONCLUSIONS AND RECOMMENDATIONS

This paper investigated PINNs for multi-step PV power generation forecasting in spatially distributed PV systems at temporal resolutions on the order of seconds under temporal data sparsity, motivated by real-time power system control requirements. A recursive forecasting framework was developed to generate high-resolution multi-step forecasts from models trained exclusively on temporally sparse data, together with a validation approach aligned with this temporal resolution mismatch.

The results demonstrate that physics-based regularization in PINNs promotes physically plausible generalization between sparse observations and mitigates error accumulation under recursive deployment. The proposed improved training strategy further enhances multi-step forecast performance by explicitly aligning model training with recursive inference, leading the improved PINN to achieve the lowest overall RMSE among all evaluated models. Although the resulting improvement in RMSE is modest, the improved PINN exhibits clear advantages under challenging conditions. In particular, it consistently outperforms persistence under rapidly evolving atmospheric conditions and demonstrates superior data efficiency compared to purely data-driven NNs in data-limited regimes. In addition, the proposed correlation-based collocation point sampling strategy improves the effectiveness of physics-based regularization relative to uniform sampling.

More broadly, the findings indicate that the principal value of physics-informed learning lies not in nominal operating conditions, but in practical regimes characterized by sparse, uncertain, or limited data, which are conditions that commonly arise in real-world power system operation.

Several directions for future work follow from this study. First, the physical model incorporated in the present PINN formulations is intentionally simplified and considers only advective cloud motion with constant wind fields over the forecast horizon. Although this assumption is appropriate for the minute-scale horizons considered here [16], future work could integrate richer physical models, including diffusive and convective processes or stochastic cloud dynamics [33], [34], to further improve the realism and effectiveness of physics-based regularization.

Second, extending the proposed PINN-based framework to model the complete forecast pipeline from cloud dynamics to PV power output would enable direct end-to-end evaluation of forecast performance. In this context, the incorporation of additional physical models of PV panel behavior, such as thermal effects and electrical response characteristics [35], represents a promising direction for further improvement.

Finally, integrating the proposed forecasting framework into closed-loop power system control applications, such as MPC for LFC, would allow a direct assessment of its impact on grid stability and operational performance.

REFERENCES

[1] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. A. Cañizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G. A. Jiménez-Estévez, and N. D. Hatziargyriou, "Trends in microgrid control," *IEEE Transactions on Smart Grid*, vol. 5, pp. 1905–1919, 2014.

[2] S. R. Sinsel, R. L. Riemke, and V. H. Hoffmann, "Challenges and solution technologies for the integration of variable renewable energy sources—a review," *Renewable Energy*, vol. 145, pp. 2271–2285, 2020.

[3] S. Impram, S. Varbak Nese, and B. Oral, "Challenges of renewable energy penetration on power system flexibility: a survey," *Energy Strategy Reviews*, vol. 31, pp. 1–12, 2020.

[4] M. Milligan, L. Bird, and D. Lew, "Integrating variable renewable energy: challenges and solutions," National Renewable Energy Laboratory, Tech. Rep. NREL/TP-6A20-60451, 2013.

[5] M. Schwenzer, M. Ay, T. Bergs, and D. Abel, "Review on model predictive control: an engineering perspective," *The International Journal of Advanced Manufacturing Technology*, vol. 117, pp. 1327–1349, 2021.

[6] A. M. Ersdal, L. Imsland, and K. Uhlen, "Model predictive load-frequency control," *IEEE Transactions on Power Systems*, vol. 31, pp. 777–785, 2016.

[7] H. H. Alhelou, M. Hamedani-Golshan, R. Zamani, E. Heydarian-Forushani, and P. Siano, "Challenges and opportunities of load frequency control in conventional, modern and future smart power systems: a comprehensive review," *Energies*, vol. 11, pp. 1–35, 2018.

[8] M. M. Amiri, M. Shadman, and S. F. Estefen, "A review of physical and numerical modeling techniques for horizontal-axis wind turbine wakes," *Renewable and Sustainable Energy Reviews*, vol. 193, pp. 1–27, 2024.

[9] T. Liu, L. Shan, M. Jiang, F. Li, F. Kong, P. Du, H. Zhu, H. H. Goh, T. A. Kurniawan, C. Huang, and D. Zhang, "Multi-dimensional data processing and intelligent forecasting technologies for renewable energy generation," *Applied Energy*, vol. 398, pp. 1–32, 2025.

[10] M. R. P. Colón, A. I. C. Salamán, D. C. Figueroa, A. Sundararajan, and M. Ferrari, "Data-driven model for photovoltaic generation: comparison with physical models using a microgrid in puerto rico," in *2024 IEEE Power & Energy Society General Meeting*, 2024, pp. 1–5.

[11] N. E. Benti, M. D. Chaka, and A. G. Semie, "Forecasting renewable energy generation with machine learning and deep learning: current advances and future prospects," *Sustainability*, vol. 15, pp. 1–33, 2023.

[12] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, pp. 422–440, 2021.

[13] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[14] A. Sorjamaa, J. Hao, N. Reyhani, Y. Ji, and A. Lendasse, "Methodology for long-term prediction of time series," *Neurocomputing*, vol. 70, pp. 2861–2869, 2007.

[15] S. Ben Taieb, G. Bontempi, A. F. Atiya, and A. Sorjamaa, "A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition," *Expert Systems with Applications*, vol. 39, pp. 7067–7083, 2012.

[16] Y. Chu, M. Li, C. F. M. Coimbra, D. Feng, and H. Wang, "Intra-hour irradiance forecasting techniques for solar power integration: a review," *iScience*, vol. 24, pp. 1–50, 2021.

[17] C. Arbizu-Barrena, J. A. Ruiz-Arias, F. J. Rodríguez-Benítez, D. Pozo-Vázquez, and J. Tovar-Pescador, "Short-term solar radiation forecasting by advecting and diffusing MSG cloud index," *Solar Energy*, vol. 155, pp. 1092–1103, 2017.

[18] T. M. Harty, W. F. Holmgren, A. T. Lorenzo, and M. Morzfeld, "Intra-hour cloud index forecasting with data assimilation," *Solar Energy*, vol. 185, pp. 270–282, 2019.

[19] N. Straub, S. Karalus, W. Herzberg, and E. Lorenz, "Satellite-based solar irradiance forecasting: replacing cloud motion vectors by deep learning," *Solar RRL*, vol. 8, pp. 1–16, 2024.

[20] J. Ranalli, E. E. M. Peerlings, and T. Schmidt, "Cloud advection and spatial variability of solar irradiance," in *2020 47th IEEE Photovoltaic Specialists Conference*, 2020, pp. 37–44.

[21] A. Mojtabi and M. O. Deville, "One-dimensional linear advection–diffusion equation: analytical and finite element solutions," *Computers & Fluids*, vol. 107, pp. 189–195, 2015.

[22] A. Baydin, B. Pearlmutter, A. Radul, and J. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of Machine Learning Research*, vol. 18, pp. 1–43, 2018.

[23] K. Cabello-Solorzano, I. Ortigosa de Araujo, M. Peña, L. Correia, and A. J. Tallón-Ballesteros, "The impact of data normalization on the accuracy of machine learning algorithms: a comparative analysis," in *18th International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2023)*, 2023, pp. 344–353.

[24] D. W. Scott, "The curse of dimensionality and dimension reduction," in *Multivariate Density Estimation: Theory, Practice and Visualization*. John Wiley & Sons, 1992, ch. 7, pp. 195–217.

[25] J. Lee Rodgers and W. A. Nicewander, "Thirteen ways to look at the correlation coefficient," *The American Statistician*, vol. 42, pp. 59–66, 1988.

[26] X. Liu, X. Zhu, P. Qiu, and W. Chen, "A correlation-matrix-based hierarchical clustering method for functional connectivity analysis," *Journal of Neuroscience Methods*, vol. 211, pp. 94–102, 2012.

[27] F. Nielsen, "Hierarchical clustering," in *Introduction to HPC with MPI for Data Science*. Springer International Publishing, 2016, pp. 195–211.

[28] Y. Chen, "A tutorial on kernel density estimation and recent advances," 2017.

[29] L. C. Evans, *Partial Differential Equations*, ser. Graduate Studies in Mathematics. American Mathematical Society, 2010, vol. 19.

[30] Solcast. Historical time series. [Online]. Available: https://solcast.com/time-series

[31] V. Tzallas, A. Hünerbein, M. Stengel, K. F. Meirink, N. Benas, J. Trentmann, and A. Macke, "CRAAS: a european cloud regime dAtAset based on the CLAAS-2.1 climate data record," *Remote Sensing*, vol. 14, pp. 1–25, 2022.

[32] R. Perez, M. David, T. Hoff, M. Jamaly, S. Kivalov, J. Kleissl, P. Lauret, and M. Perez, "Spatial and temporal variability of solar energy," *Foundations and Trends® in Renewable Energy*, vol. 1, pp. 1–44, 2016.

[33] J. Ranalli and E. Peerlings, "Cloud advection model of solar irradiance smoothing by spatial aggregation," *Journal of Renewable and Sustainable Energy*, vol. 13, pp. 1–13, 2021.

[34] T. N. Palmer and P. D. Williams, "Introduction. Stochastic physics and climate modelling," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 366, pp. 2419–2425, 2008.

[35] K. Wang, L. Wang, Q. Meng, C. Yang, Y. Lin, J. Zhu, Z. Zhao, C. Zhou, C. Zheng, and X. Gao, "Accurate photovoltaic power prediction via temperature correction with physics-informed neural networks," *Energy*, vol. 328, pp. 1–14, 2025.

Bibliography

[1] X. G. Agoua, R. Girard, and G. Kariniotakis. "Short-term spatio-temporal forecasting of photovoltaic power production". In: *IEEE Transactions on Sustainable Energy* 9 (2018), pp. 538–546.

[2] A. Ahmed and M. Khalid. "A review on the selected applications of forecasting models in renewable power systems". In: *Renewable and Sustainable Energy Reviews* 100 (2019), pp. 9–21.

[3] M. Ajith and M. Martínez-Ramón. "Deep learning based solar radiation micro forecast by fusion of infrared cloud images and radiation data". In: *Applied Energy* 294 (2021), pp. 1–14.

[4] M. S. Alam et al. "Solar and wind energy integrated system frequency control: a critical review on recent developments". In: *Energies* 16 (2023), pp. 1–31.

[5] H. H. Alhelou et al. "Challenges and opportunities of load frequency control in conventional, modern and future smart power systems: a comprehensive review". In: *Energies* 11 (2018), pp. 1–35.

[6] M. P. Almeida et al. "Comparative study of PV power forecast using parametric and nonparametric PV models". In: *Solar Energy* 155 (2017), pp. 854–866.

[7] M. M. Amiri, M. Shadman, and S. F. Estefen. "A review of physical and numerical modeling techniques for horizontal-axis wind turbine wakes". In: *Renewable and Sustainable Energy Reviews* 193 (2024), pp. 1–27.

[8] C. Arbizu-Barrena et al. "Short-term solar radiation forecasting by advecting and diffusing MSG cloud index". In: *Solar Energy* 155 (2017), pp. 1092–1103.

[9] E. Arslan Tuncar, Ş. Sağlam, and B. Oral. "A review of short-term wind power generation forecasting methods in recent technological trends". In: *Energy Reports* 12 (2024), pp. 197–209.

[10] F. Barbieri, S. Rajakaruna, and A. Ghosh. "Very short-term photovoltaic power forecasting with cloud modeling: a review". In: *Renewable and Sustainable Energy Reviews* 75 (2017), pp. 242–263.

[11] A. Baydin et al. “Automatic differentiation in machine learning: a survey”. In: *Journal of Machine Learning Research* 18 (2018), pp. 1–43.

[12] S. Ben Taieb et al. “A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition”. In: *Expert Systems with Applications* 39 (2012), pp. 7067–7083.

[13] S. Bengio et al. *Scheduled sampling for sequence prediction with recurrent neural networks*. 2015. DOI: 10.48550/arXiv.1506.03099.

[14] N. E. Benti, M. D. Chaka, and A. G. Semie. “Forecasting renewable energy generation with machine learning and deep learning: current advances and future prospects”. In: *Sustainability* 15 (2023), pp. 1–33.

[15] G. Bontempi and S. Ben Taieb. “Conditionally dependent strategies for multiple-step-ahead prediction in local learning”. In: *International Journal of Forecasting* 27 (2011), pp. 689–699.

[16] R. B. Buenestado. *The Iberian peninsula blackout — causes, consequences, and challenges ahead*. Tech. rep. Rice University’s Baker Institute for Public Policy, 2025.

[17] K. Cabello-Solorzano et al. “The impact of data normalization on the accuracy of machine learning algorithms: a comparative analysis”. In: *18th International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2023)*. 2023, pp. 344–353.

[18] A. Caradot et al. *Provably accurate adaptive sampling for collocation points in physics-informed neural networks*. 2025. DOI: 10.48550/arXiv.2504.00910.

[19] S. Chai et al. “A robust spatiotemporal forecasting framework for photovoltaic generation”. In: *IEEE Transactions on Smart Grid* 11 (2020), pp. 5370–5382.

[20] Y. Chen. *A tutorial on kernel density estimation and recent advances*. 2017. DOI: 10.48550/arXiv.1704.03924.

[21] Z. Chen, Y. Liu, and H. Sun. “Physics-informed learning of governing equations from scarce data”. In: *Nature Communications* 12 (2021), pp. 1–13.

[22] Y. Chu et al. “Intra-hour irradiance forecasting techniques for solar power integration: a review”. In: *iScience* 24 (2021), pp. 1–50.

[23] M. R. P. Colón et al. “Data-driven model for photovoltaic generation: comparison with physical models using a microgrid in Puerto Rico”. In: *2024 IEEE Power & Energy Society General Meeting*. 2024, pp. 1–5.

[24] S. Cuomo et al. “Scientific machine learning through physics–informed neural networks: where we are and what’s next”. In: *Journal of Scientific Computing* 92 (2022), pp. 1–62.

[25] P. Dogoulis, K. Tit, and M. Cordy. *KCLNet: physics-informed power flow prediction via constraints projections*. 2025. DOI: 10.48550/arXiv.2506.12902.

[26] A. Dolara, S. Leva, and G. Manzolini. “Comparison of different physical models for PV power output prediction”. In: *Solar Energy* 119 (2015), pp. 83–99.

[27] A. M. Ersdal, L. Imsland, and K. Uhlen. “Model predictive load-frequency control”. In: *IEEE Transactions on Power Systems* 31 (2016), pp. 777–785.

[28] European Environment Agency. *Share of energy consumption from renewable sources in Europe*. 2025. URL: https://www.eea.europa.eu/en/analysis/indicators/share-of-energy-consumption-from.

[29] L. C. Evans. *Partial Differential Equations*. Vol. 19. Graduate Studies in Mathematics. American Mathematical Society, 2010.

[30] S. A. Faroughi et al. "Physics-guided, physics-informed, and physics-encoded neural networks and operators in scientific computing: fluid and solid mechanics". In: *Journal of Computing and Information Science in Engineering* 24 (2024), pp. 1–31.

[31] J. Florido et al. "Investigating guiding information for adaptive collocation point sampling in PINNs". In: *Computational Science – ICCS 2024*. Ed. by L. Franco et al. Springer Nature Switzerland, 2024, pp. 323–337.

[32] M. M. Fouad, L. A. Shihata, and E. I. Morgan. "An integrated review of factors influencing the performance of photovoltaic panels". In: *Renewable and Sustainable Energy Reviews* 80 (2017), pp. 1499–1511.

[33] P. Gupta and R. Singh. "PV power forecasting based on data-driven models: a review". In: *International Journal of Sustainable Engineering* 14 (2021), pp. 1733–1755.

[34] Z. Hao et al. *Physics-informed machine learning: a survey on problems, methods and applications*. 2023. DOI: 10.48550/arXiv.2211.08064.

[35] T. M. Harty et al. "Intra-hour cloud index forecasting with data assimilation". In: *Solar Energy* 185 (2019), pp. 270–282.

[36] S. Homan, N. Mac Dowell, and S. Brown. "Grid frequency volatility in future low inertia scenarios: challenges and mitigation options". In: *Applied Energy* 290 (2021), pp. 1–24.

[37] Z. Hu et al. "Improved multistep ahead photovoltaic power prediction model based on LSTM and self-attention with weather forecast data". In: *Applied Energy* 359 (2024), pp. 1–15.

[38] B. Huang and J. Wang. "Applications of physics-informed neural networks in power systems - a review". In: *IEEE Transactions on Power Systems* 38 (2023), pp. 572–588.

[39] S. Impram, S.. Varbak Nese, and B. Oral. "Challenges of renewable energy penetration on power system flexibility: a survey". In: *Energy Strategy Reviews* 31 (2020), pp. 1–12.

[40] S. Inage. "Development of an advection model for solar forecasting based on ground data first report: development and verification of a fundamental model". In: *Solar Energy* 153 (2017), pp. 414–434.

[41] P. Ineichen and R. Perez. "A new airmass independent formulation for the linke turbidity coefficient". In: *Solar Energy* 73 (2002), pp. 151–157.

[42] G. E. Karniadakis et al. "Physics-informed machine learning". In: *Nature Reviews Physics* 3 (2021), pp. 422–440.

[43] K. Kashinath et al. "Physics-informed machine learning: case studies for weather and climate modelling". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379 (2021), pp. 1–36.

[44] D. Kline. "Methods for multi-step time series forecasting with neural networks". In: *Neural Networks in Business Forecasting*. 2004, pp. 226–250.

[45] A. Krishnapriyan et al. “Characterizing possible failure modes in physics-informed neural networks”. In: *Advances in Neural Information Processing Systems*. Vol. 34. 2021, pp. 26548–26560.

[46] J. Kruse et al. “Physics-informed machine learning for power grid frequency modeling”. In: *PRX Energy* 2 (2023), pp. 1–20.

[47] J. Lee Rodgers and W. A. Nicewander. “Thirteen ways to look at the correlation coefficient”. In: *The American Statistician* 42 (1988), pp. 59–66.

[48] X. Liang. “Emerging power quality challenges due to integration of renewable energy sources”. In: *IEEE Transactions on Industry Applications* 53 (2017), pp. 855–866.

[49] W. Lin et al. “Multi-step prediction of photovoltaic power based on two-stage decomposition and BILSTM”. In: *Neurocomputing* 504 (2022), pp. 56–67.

[50] T. Liu et al. “Multi-dimensional data processing and intelligent forecasting technologies for renewable energy generation”. In: *Applied Energy* 398 (2025), pp. 1–32.

[51] X. Liu et al. “A correlation-matrix-based hierarchical clustering method for functional connectivity analysis”. In: *Journal of Neuroscience Methods* 211 (2012), pp. 94–102.

[52] G. Lu et al. “A survey of deep learning for time series forecasting: theories, datasets, and state-of-the-art techniques”. In: *Computers, Materials & Continua* 85 (2025), pp. 2403–2441.

[53] D. Markovics and M. J. Mayer. “Comparison of machine learning methods for photovoltaic power forecasting based on numerical weather prediction”. In: *Renewable and Sustainable Energy Reviews* 161 (2022), pp. 1–17.

[54] A. Mellit, A. Pavan Massi, and V. Lughi. “Deep learning neural networks for short-term photovoltaic power forecasting”. In: *Renewable Energy* 172 (2021), pp. 276–288.

[55] M. Milligan, L. Bird, and D. Lew. *Integrating variable renewable energy: challenges and solutions*. Tech. rep. NREL/TP-6A20-60451. National Renewable Energy Laboratory, 2013, p. 14.

[56] A. Mojtabi and M. O. Deville. “One-dimensional linear advection–diffusion equation: analytical and finite element solutions”. In: *Computers & Fluids* 107 (2015), pp. 189–195.

[57] K. Moslehi and R. Kumar. “A reliability perspective of the smart grid”. In: *IEEE Transactions on Smart Grid* 1 (2010), pp. 57–64.

[58] M. A. Nabian, R. J. Gladstone, and H. Meidani. “Efficient training of physics-informed neural networks via importance sampling”. In: *Computer-Aided Civil and Infrastructure Engineering* 36 (2021), pp. 962–977.

[59] M. Nagda et al. *PIANO: physics informed autoregressive network*. 2025. DOI: 10.48550/arXiv.2508.16235.

[60] F. Nielsen. “Hierarchical clustering”. In: *Introduction to HPC with MPI for Data Science*. Springer International Publishing, 2016, pp. 195–211.

[61] D. E. Olivares et al. “Trends in microgrid control”. In: *IEEE Transactions on Smart Grid* 5 (2014), pp. 1905–1919.

[62] T. N. Palmer and P. D. Williams. "Introduction. Stochastic physics and climate modelling". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 366 (2008), pp. 2419–2425.

[63] H. Panamtash et al. "Very short-term solar power forecasting using a frequency incorporated deep learning model". In: *IEEE Open Access Journal of Power and Energy* 10 (2023), pp. 517–527.

[64] A. Parisio, E. Rikos, and L. Glielmo. "A model predictive control approach to microgrid operation optimization". In: *IEEE Transactions on Control Systems Technology* 22 (2014), pp. 1813–1827.

[65] K. Park, J. Kim, and J. Seo. *PINT: physics-informed neural time series models with applications to long-term inference on WeatherBench 2m-temperature data*. 2025. DOI: 10.48550/arXiv.2502.04018.

[66] R. Perez et al. "Spatial and temporal variability of solar energy". In: *Foundations and Trends® in Renewable Energy* 1 (2016), pp. 1–44.

[67] D. V. Pombo et al. "Benchmarking physics-informed machine learning-based short term PV-power forecasting tools". In: *Energy Reports* 8 (2022), pp. 6512–6520.

[68] M. Raissi, P. Perdikaris, and G. E. Karniadakis. "Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: *Journal of Computational Physics* 378 (2019), pp. 686–707.

[69] J Ranalli, E. E. M. Peerlings, and T. Schmidt. "Cloud advection and spatial variability of solar irradiance". In: *2020 47th IEEE Photovoltaic Specialists Conference*. 2020, pp. 37–44.

[70] J. Ranalli and E. Peerlings. "Cloud advection model of solar irradiance smoothing by spatial aggregation". In: *Journal of Renewable and Sustainable Energy* 13 (2021), pp. 1–13.

[71] A. Sabadus et al. "A cross-sectional survey of deterministic PV power forecasting: progress and limitations in current approaches". In: *Renewable Energy* 226 (2024), pp. 1–22.

[72] M. Schwenzer et al. "Review on model predictive control: an engineering perspective". In: *The International Journal of Advanced Manufacturing Technology* 117 (2021), pp. 1327–1349.

[73] D. W. Scott. "The curse of dimensionality and dimension reduction". In: *Multivariate Density Estimation: Theory, Practice and Visualization*. John Wiley & Sons, 1992. Chap. 7, pp. 195–217.

[74] M. Sengupta and A. Andreas. *Oahu Solar Measurement Grid (1-Year Archive): 1-Second Solar Irradiance; Oahu, Hawaii (data)*. 2014. DOI: 10.7799/1052451.

[75] P. Sharma et al. "A review of physics-informed machine learning in fluid mechanics". In: *Energies* 16 (2023), pp. 1–21.

[76] J. Simeunović et al. "Interpretable temporal-spatial graph attention network for multi-site PV power forecasting". In: *Applied Energy* 327 (2022), pp. 1–15.

[77] S. R. Sinsel, R. L. Riemke, and V. H. Hoffmann. "Challenges and solution technologies for the integration of variable renewable energy sources—a review". In: *Renewable Energy* 145 (2020), pp. 2271–2285.

[78] J. P. Snyder. *Map projections: a working manual*. Tech. rep. 1395. U.S. Government Printing Office, 1987, p. 385.

[79] S. Sobri, S. Koohi-Kamali, and N. A. Rahim. "Solar photovoltaic generation forecasting methods: a review". In: *Energy Conversion and Management* 156 (2018), pp. 459–497.

[80] Solcast. *Historical Time Series*. URL: https://solcast.com/time-series.

[81] A. Sorjamaa and A. Lendasse. "Time series prediction using DirRec strategy". In: *The European Symposium on Artificial Neural Networks*. 2006, pp. 143–148.

[82] A. Sorjamaa et al. "Methodology for long-term prediction of time series". In: *Neurocomputing* 70 (2007), pp. 2861–2869.

[83] A. Staniforth and J. Côté. "Semi-lagrangian integration schemes for atmospheric models—a review". In: *Monthly Weather Review* 119 (1991), pp. 2206–2223.

[84] N. Straub et al. "Satellite-based solar irradiance forecasting: replacing cloud motion vectors by deep learning". In: *Solar RRL* 8 (2024), pp. 1–16.

[85] S. B. Taieb and A. F. Atiya. "A bias and variance analysis for multistep-ahead time series forecasting". In: *IEEE Transactions on Neural Networks and Learning Systems* 27 (2016), pp. 62–76.

[86] V. Tzallas et al. "CRAAS: a European cloud regime dAtAset based on the CLAAS-2.1 climate data record". In: *Remote Sensing* 14 (2022), pp. 1–25.

[87] H. Wang et al. "Solar irradiance forecasting based on direct explainable neural network". In: *Energy Conversion and Management* 226 (2020), pp. 1–11.

[88] K. Wang et al. "Accurate photovoltaic power prediction via temperature correction with physics-informed neural networks". In: *Energy* 328 (2025), pp. 1–14.

[89] C. Wu et al. "A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks". In: *Computer Methods in Applied Mechanics and Engineering* 403 (2023), pp. 1–29.

[90] D. Yang et al. "Very short term irradiance forecasting using the lasso". In: *Solar Energy* 114 (2015), pp. 314–326.

[91] W. Zhang et al. *Physics informed neural networks (PINNs) as intelligent computing technique for solving partial differential equations: limitation and future prospects*. 2024. DOI: 10.48550/arXiv.2411.18240.

[92] X. Zhong et al. "Explainable machine learning in materials science". In: *npj Computational Materials* 8 (2022), pp. 1–19.

Glossary

KDE	kernel density estimation
LFC	load frequency control
MAE	mean absolute error
ML	machine learning
MPC	model predictive control
MSE	mean squared error
NN	neural network
NREL	National Laboratory of the Rockies
NWP	numerical weather prediction
PIML	physics-informed machine learning
PINN	physics-informed neural network
PV	photovoltaic
RES	renewable energy source
RMSE	root mean squared error
UTC	Coordinated Universal Time
VE	variability error