energies

MDPI

Article

MeteoSynth: A Python Package for Non-Parametric Stochastic Weather Series Generation in Renewable Energy and Building Performance Simulations

Nikolaos C. Papadakis

Mechanical Engineering Department, Polytechnic School, Hellenic Mediterranean University, Estavromenos, 71410 Heraklion, Greece; npap@hmu.gr; Tel.: +30-2810379722

Abstract

Evaluating renewable energy system performance and building energy demand typically relies on single-year historical profiles or Typical Meteorological Years (TMYs). However, deterministic weather series fail to capture inter-annual variability, extreme operational events, and provide distribution information necessary for performance assessment, control optimization, and Global Sensitivity Analysis (GSA). To address these limitations, this paper introduces MeteoSynth, an open-source Python package designed for high-throughput, non-parametric generation of physically consistent synthetic time series for individual environmental variables. The package is variable-agnostic: any continuous hourly series can be modeled once its physical support and units are declared. This paper, like the package defaults, focuses on solar irradiance, air temperature, and wind speed. These three were chosen mainly because their physical constraints differ: temperature may take negative values, wind speed is non-negative, and irradiance is non-negative and also collapses to a degenerate zero state every night. Together they show how the generator handles each of these cases in general, and the limitations that remain are also examined. A second and practical reason is that the European Commission's Photovoltaic Geographical Information System (PVGIS) supplies all three for any location. The present release is a parallel univariate generator: each variable is simulated independently, so the variables share the diurnal and seasonal cycle but not their day-to-day departures from it. Joint multivariate generation is identified as the main next development step. MeteoSynth models diurnal weather dynamics using time-inhomogeneous first-order Markov chains governed by two-dimensional Kernel Density Estimation (2D KDE). To ensure physical validity and calendar continuity, the framework incorporates features such as prevention of non-physical negative values with minimal domain distortion, a robust degenerate state classification engine to handle zero-variance night-time solar irradiance, a climatological circular rolling-window calendar (366-day calendar) that eliminates New Year boundary jumps while accommodating leap years, and date-aligned cross-midnight transition models ensuring multi-day series continuity. It additionally offers optional integration with the PVGIS database, a user-friendly API supporting both rapid monthly subsetting and high-accuracy rolling-window training, and an optimized ensemble engine capable of generating thousands of synthetic weather realizations in a single calendar walk, such that generating a single synthetic year and large ensembles of 1000 to 100,000 years takes virtually the same amount of execution time due to internal model-reuse optimizations.

Keywords: synthetic weather generation; non-parametric Markov chains; building energy simulation; weather stochasticity; open-source Python software

Check for updates

Academic Editor: Xiaohu Yang

Received: 24 October 2026
Revised: 20 November 2026
Accepted: 22 November 2026
Published: 24 November 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Energy system design, building thermal simulation, and renewable power integration rely fundamentally on meteorological time-series inputs such as global solar irradiance (G), outdoor air temperature (T_{air}), and wind speed (v_{wind}). Historically, engineering practice and building performance codes have relied on deterministic weather datasets, most notably single historical observations or synthesized Typical Meteorological Years (TMYs). While TMYs provide a standardized profile for average energy consumption estimates, they are constructed specifically to represent median climatic conditions over multi-decade records. By smoothing out natural climate variability and suppressing extreme meteorological fluctuations, TMY profiles fail to expose energy systems to the tail risks, multi-day weather persistent events, and operational stresses that dictate system reliability, thermal comfort violations, battery storage autonomy, and financial performance.

To quantify operational risks and optimize energy system resilience, modern engineering workflows increasingly require Monte Carlo simulation frameworks driven by synthetic weather ensembles. Generating synthetic meteorological series requires balancing physical realism, computational scalability, and statistical fidelity. Existing parametric weather generators often impose strict statistical distribution assumptions (e.g., Gaussian or Weibull distributions) that fail to capture the skewed, multi-modal, and time-varying joint probability distributions characteristic of hourly microclimates. Furthermore, simple autoregressive or discrete transition matrix models frequently introduce boundary discontinuities across days, exhibit non-physical out-of-bounds artifacts (e.g., negative solar irradiance or negative wind speeds), or suffer from severe memory bottlenecks when scaled to large Monte Carlo ensembles.

In this work, we present MeteoSynth, an open-source Python library designed to generate statistically faithful, physically consistent, hourly synthetic weather series and ensembles from historical meteorological measurements or optionally from automated web API data services. MeteoSynth adopts a non-parametric two-dimensional Kernel Density Estimation (2D KDE) first-order Markov chain framework. By fitting time-inhomogeneous transition models across hourly steps, MeteoSynth captures non-linear diurnal dynamics without relying on parametric distribution assumptions. Nothing in this formulation is specific to a particular meteorological quantity. The generator treats any continuous hourly series in the same way, and each variable enters only through a short declaration of its physical support and units. Relative humidity, dew-point temperature or diffuse irradiance, for example, can therefore be added without changing the package. In its present release, MeteoSynth is a parallel univariate generator: each variable is fitted and simulated on its own. The variables share the calendar and the hour of day, so their common diurnal and seasonal cycle is preserved, but the dependence between their departures from that cycle (for example, that a cloudy winter day in Crete also tends to be cooler and windier) is not. The observed correlations that are therefore not reproduced are reported in Section 4.4.5, and joint multivariate generation is the main next step of the development roadmap (Section 6.3).

The package introduces several key architectural and mathematical innovations designed for physical consistency and computational efficiency. First, a physical support truncation mechanism strictly clamps values within realistic environmental bounds (e.g., enabling guarantees for non-negative solar irradiance and wind speed, i.e., $G \geq 0$, $v_{wind} \geq 0$) without introducing artificial distributional bias. To address invariant intervals such as nocturnal solar radiation, a robust degenerate state classification handles zero-variance transition periods, avoiding singularities in kernel density covariance matrices. Continuous seasonal transitions are maintained using an all-leap 366-day circular rolling window ($\pm$n days) that treats February 29 as a persistent calendar slot and wraps smoothly across the New Year boundary. In parallel, a dedicated transition model conditioned on consecutive

calendar dates preserves date-aligned cross-midnight continuity, preventing unphysical step artifacts across multi-day sequences. Finally, the generator employs a single-calendar-walk ensemble architecture that evaluates precomputed conditional quantile grids, making the generation of thousands of synthetic realizations approximately two orders of magnitude faster than generating them one by one.

MeteoSynth provides an intuitive high-level API, including direct integration with the Photovoltaic Geographical Information System (PVGIS) [1] via pvlib [2–4]. PVGIS is a free web service of the European Commission's Joint Research Centre. Besides photovoltaic yield calculations, it serves multi-year hourly series of solar irradiance, air temperature, and wind speed for any location. Its irradiance over Europe is derived from satellite observations (the PVGIS-SARAH3 database for the site of Section 4), while its air temperature and wind speed come from the ERA5 reanalysis [5]. The PVGIS connector is optional: the generator itself accepts any hourly record, such as ground-station measurements or reanalysis data obtained elsewhere. MeteoSynth also offers flexible model training options (fast monthly subsetting vs. high-accuracy daily rolling windows). The paper focuses on three variables, solar irradiance, air temperature, and wind speed, which are also the package defaults. They were selected primarily because their physical constraints differ, so that together they exercise every safeguard of the generator. Air temperature is bounded only by absolute zero and routinely takes negative values. Wind speed is non-negative but rarely reaches its lower bound. Irradiance is non-negative, sits on that bound at dawn and dusk, and is exactly zero throughout the night, which makes it degenerate. A second and practical reason is that PVGIS provides all three for any location, which makes the case study easy to reproduce. In this paper, we detail the mathematical foundations, software design, empirical calibration metrics, and practical deployment of MeteoSynth across four representative energy engineering domains: building energy Monte Carlo simulation, stochastic HVAC control strategy optimization, PV energy yield risk quantification, and Global Sensitivity Analysis (GSA).

2. Literature Review

2.1. The Simulation Paradigm: Climatological Synthesis vs. Operational Forecasting

2.1.1. Operational Forecasting vs. Stochastic Climatological Simulation

The importance of the uncertainty of weather conditions on building energy assessment has received increasing attention in recent years [6]. Operational forecasting focuses on estimating immediate atmospheric trajectories over discrete, short-term horizons ranging from intraday timesteps 15 min to 6 h) to day-ahead periods 24 to 72 h. Conditioned on real-time sensor telemetry, initial atmospheric state vectors x_0, and Numerical Weather Prediction (NWP) physics runs, forecasting models aim to minimize point-wise deviation metrics, such as Root Mean Square Error (RMSE) and Mean Absolute Error (MAE), or to calibrate prediction intervals via quantile loss and Continuous Ranked Probability Scores [7,8]. The primary operational objective is deterministic scheduling [9,10], spinning reserve allocation [11,12], and short-term grid balancing [13].

In contrast, long-term renewable integration, infrastructure investment, and system adequacy planning require an entirely distinct modeling paradigm: unconditional stochastic climatological simulation [14–16]. Climatological synthesis spans continuous annual (8760 h) or multi-year/decade horizons [15,17]. Rather than forecasting exact realization paths at deterministic future calendar timestamps, its mathematical objective is sampling from an underlying joint probability distribution that preserves long-term climatological characteristics. These include diurnal and seasonal periodicity, multiscale autocorrelation decay, non-Gaussian marginal distributions, and non-linear cross-variable dependencies across solar irradiance (GHI,DNI), wind velocity (v), and ambient temperature (T) [18–20].

Short-horizon forecasting and long-term climatological synthesis are similar but distinct concepts, and confusing them leads to severe misspecifications in energy system models' performance assessment [21]. While forecasting architectures are optimized to minimize local variance around expected conditional means, energy planning models are particularly sensitive to distributional tails, compound multi-day energy deficits, and inter-annual volatility, features systematically suppressed by point-forecast objectives.

2.1.2. Historical Evolution and Structural Limitations of Typical Meteorological Years (TMYs)

To manage the computational complexity of multi-decadal climate datasets in engineering design, early modeling efforts introduced standardized, single-year representative profiles, such as Typical Meteorological Years (TMY, TMY2, TMY3), International Weather for Energy Calculations (IWEC), and ISO 15927-4 standards [22,23]. Originally developed for commercial building heating and cooling simulations, the TMY algorithm uses the Finkelstein–Schafer (FS) non-parametric test to evaluate 30-year historical station records [24,25]. It selects the candidate month whose empirical cumulative distribution function (ECDF) closest matches the long-term composite across key meteorological parameters, concatenating twelve representative months into an artificial annual sequence of 8760 h.

Although TMYs remain widely embedded in engineering software, Pfenninger and Staffell [26,27] identify a set of fundamental structural deficiencies that emerge when such datasets are applied to weather-dependent power system modeling. Because the FS metric explicitly rewards convergence toward median distributions, anomalous atmospheric conditions are systematically excluded. The independent concatenation of twelve distinct months compounds this by introducing step-discontinuities across monthly transition boundaries, breaking multi-day thermal and wind speed persistence. More fundamentally, the method reduces multi-year stochastic climate regimes into a single deterministic trajectory, which prevents Monte Carlo risk assessment and hides inter-annual resource volatility. Traditional TMY datasets further assume a stationary climate, ignoring structural shifts in ambient temperature distributions and cloud-cover patterns driven by anthropogenic climate change.

2.1.3. Meteorological Uncertainty and Multi-Year Energy System Resilience

Modern decarbonized power systems rely heavily on variable renewable energy (VRE), making grid stability directly dependent on compound atmospheric states [28]. A critical vulnerability is extreme events; for example, the Dunkelflaute (dark doldrums), which are events with little direct sunlight and low wind speed, are critical for sizing long-duration energy storage (e.g., green hydrogen, pumped hydro, flow batteries) and ensuring resource adequacy [29].

Evaluating metrics such as Loss of Load Expectation (LOLE) [30], Expected Energy Not Served (EENS) [31,32], and capacity credit [33,34] requires large ensembles of multi-year weather profiles that reflect true joint probabilities [16]. The paradigm of synthetic time series generation (STSG) resolves the trade-off between the computational burden of processing raw multi-decadal gridded reanalysis products (e.g., ERA5 [5], MERRA-2 [35]) and the structural oversimplifications of static TMYs. STSG method ought to provide an automated mechanism to generate unbounded, statistically sound, and physically constrained meteorological scenarios for probabilistic energy system design.

2.2. Paradigm I: Classical Parametric and Stochastic Methods

Classical parametric and analytical regression methods have historically served, and continue to serve, as the foundational workhorse for synthetic weather generation. Estab-

lished software suites and early engineering guidelines routinely adopt empirical regressions, clear-sky radiation formulations [36], and low-order harmonic Fourier series [37] to isolate the deterministic diurnal and seasonal trajectories of ambient temperature and solar insolation. Because these analytical regression frameworks isolate the predictable astronomical and cyclical drivers of local microclimates, they remain the standard industry baseline for establishing long-term climate normals and deterministic baselines prior to stochastic perturbation.

Despite their distinct functional expressions across temperature, wind, and irradiance, these classical approaches are unified by a single core theoretical principle: non-stationary environmental time series can be formally decomposed into an analytical deterministic trend and a stationary stochastic residual process [38–40]. By projecting complex atmospheric dynamics onto well-defined parametric probability distributions and linear time-invariant operators, Paradigm I methods provide mathematical transparency, computational tractability, and rigorous parameter interpretability. The theoretical formulations governing these parameter-specific decompositions and their coupled multivariate integration are developed in the sections that follow.

2.2.1. Solar Irradiance Parameter-Specific Formulations

Because raw Global Horizontal Irradiance (GHI) and Direct Normal Irradiance (DNI) exhibit strong deterministic diurnal profiles bounded strictly by zero at night, direct autoregressive modeling is mathematically invalid. Classical frameworks resolve this by decomposing irradiance through a deterministic clear-sky model ($GHI_{cs,t}$), yielding a normalized, dimensionless stochastic metric known as the clearness index:

$$k_t = \frac{GHI_t}{GHI_{cs,t}}$$

To model the stochastic dynamics of k_t, classical approaches deploy discrete-state Markov Transition Matrices (MTMs) or continuous autoregressive formulations [41–43]. In Markovian schemes, k_t is discretized into discrete cloud-cover states (e.g., clear, overcast, partially cloudy), where transition probabilities between states at time t and $t + 1$ are calibrated from historical transition frequencies [44]. Alternatively, in continuous models such as the Time-dependent Autoregressive Gaussian (TAG) framework, k_t is transformed into a standardized Gaussian latent variable via empirical cumulative distribution function (ECDF) inversion or logit transforms, and subsequently modeled using a low-order autoregressive process, AR(p), to capture cloud-cover persistence [42].

2.2.2. Wind Speed Parameter-Specific Formulations

Unlike irradiance, surface wind velocity (v_t) is a continuous, strictly non-negative variable characterized by right-skewed, non-Gaussian marginal distributions [38]. Classical stochastic synthesis approaches wind through a two-step transformation:

1. Marginal distribution fitting using two-parameter Weibull or Rayleigh distributions:

$$f(v;k,c) = \frac{k}{c}\left(\frac{v}{c}\right)^{k-1}\exp\left(-\left(\frac{v}{c}\right)^{k}\right)$$

2. Probability Integral Transform (PIT) mapping to project empirical wind speeds into standard Gaussian white-noise spaces, where linear Autoregressive Moving Average (ARMA) or fractionally integrated FARIMA(p,d,q) models are fitted [38,45].

FARIMA processes are used to capture the long-memory persistence and Hurst phenomena typical of synoptic-scale atmospheric pressure fields [45]. In continuous time,

mean-reverting Ornstein–Uhlenbeck (OU) stochastic differential equations driven by Lévy jump processes are implemented to simulate sudden wind gusts and ramp events:

$$dv_t = \theta(\mu(t) - v_t)dt + \sigma(t)dW_t + dJ_t$$

where θ governs the speed of mean reversion, $\mu(t)$ tracks the diurnal mean trajectory.

Pesch et al. [46] used a Markov-chain-related statistical approach and introduced a variable second lag in order to improve the reproduction of autocorrelation characteristics in wind time series, which was lacking in earlier attempts [47–50].

2.2.3. Ambient Temperature Parameter-Specific Formulations

Ambient dry-bulb temperature (T_t) exhibits strong multiscale periodicity driven by annual insolation cycles and diurnal surface flux. Classical models represent temperature as the superposition of a low-order truncated Fourier series and a mean-reverting autoregressive error term:

$$T_t = S(t) + X_t,\ S(t) = A_0 + \sum_{m=1}^{M}\left(A_m\cos\left(\frac{2\pi mt}{365.25}\right) + B_m\sin\left(\frac{2\pi mt}{365.25}\right)\right)$$

The residual anomaly X_t is typically modeled as a stationary AR(1) or AR(2) process [51]. To capture seasonal shifts in temperature variability, the variance of the innovation noise $\varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$ is modulated using seasonal heteroskedasticity formulations (e.g., Fourier-parameterized standard deviations or Generalized Autoregressive Conditional Heteroskedasticity GARCH models).

2.2.4. Coupled Multi-Variable Weather Generators

To generate physically coherent meteorological vectors $\mathbf{X}_t = [GHI_t, v_t, T_t]^\top$, single-variable formulations can be coupled through cross-variable correlation structures. The standard framework in the classical literature is the Richardson-type Weather Generator (WGEN) and its derivatives (e.g., CLIGEN, Met-Nomad, LARS-WG) [37,52,53]. In the classic WGEN architecture, precipitation occurrence serves as the primary governing Markovian state (wet/dry days). Because Richardson-type generators were historically formulated at a daily resolution for agricultural and hydrological applications, synthetic time series generation for energy system analysis requires sub-daily downscaling [53]. This is typically achieved via Vector Autoregressive (VAR) models or deterministic diurnal shaping functions that disaggregate daily totals (such as total daily global radiation and daily temperature extremes [$T_{\min}$,$T_{\max}$]) into continuous hourly or sub-hourly trajectories [42].

2.2.5. Structural Assumptions and Microclimate Breakdown

Classical parametric generators are efficient to run and produce continuous trajectories of arbitrary length; however, they exhibit fundamental structural limitations when applied to modern renewable energy system modeling. Richardson-style formulations represent the association between variables through a single linear correlation matrix applied to Gaussian residuals, which restricts the model in two ways. Relationships between variables are forced to be linear; consequently, they fail to reproduce non-linear inter-variable dynamics, such as the threshold-driven decoupling of temperature and wind speed during atmospheric inversions. More importantly, Gaussian dependence drives extremes towards independence: the more adverse the conditions, the less likely the model is to generate them jointly. Compound low-production events are therefore systematically underrepresented. Dunkelflaute conditions, in which sustained low irradiance coincides with wind speeds below turbine cut-in [54], arise from a single synoptic cause and are often accompanied by a cold anomaly that raises demand as supply collapses [28,29]. A covariance matrix fitted to the full record reproduces the average association between irradiance,

wind speed and air temperature while thinning precisely the region where all three are simultaneously adverse.

A further limitation concerns the presumption of one homogeneous generating mechanism. Fitting a single analytical family in complex terrains—a unimodal Weibull to wind speed, a Gaussian to clearness-index residuals—violates this; e.g., coastal sites alternate between land- and sea-breeze regimes, valley sites between drainage and synoptic flow, producing multimodal distributions better described as switching between regimes than fluctuating about a single mean. Parametric models can handle this, as Markov-switching formulations demonstrate [40], but only by fixing the number of regimes in advance and re-estimating them site by site, which forfeits the transferability that made the paradigm attractive.

2.3. Paradigm II: Non-Parametric, Empirical, and Copula-Based Approaches

To overcome the restrictive distributional assumptions and linear coupling constraints of classical parametric models, non-parametric and empirical synthesis techniques eliminate predefined functional forms for marginal probability distributions and cross-variable dependencies. By sampling directly from empirical cumulative distribution functions (ECDFs), historical analog states, or flexible copula dependence structures, this paradigm preserves multimodal features, skewness, and non-linear joint interactions across solar irradiance (GHI,DNI), wind velocity (v), and ambient temperature (T).

2.3.1. Distribution-Free Resampling and Density Estimation

1. ***k***-Nearest Neighbor (***k***-NN) Local Bootstrapping:

Originally formulated for hydrometeorological simulation, k-NN resampling [55,56] constructs synthetic sequences by identifying historical atmospheric analogs conditioned on recent meteorological states [57–59]. At simulation time step t, a feature vector $\mathbf{z}_t \in \mathbb{R}^d$ is constructed from current or lagging meteorological metrics (e.g., daily mean clearness index, mean wind speed, temperature anomalies). The historical database is queried within a moving seasonal window (typically ±7 to ±15 days across all historical years) to identify the k nearest historical analog days or hours using a distance metric, such as the Mahalanobis or weighted Euclidean distance:

$$D(\mathbf{z}_t, \mathbf{z}_i) = \sqrt{(\mathbf{z}_t - \mathbf{z}_i)^\top \mathbf{W} (\mathbf{z}_t - \mathbf{z}_i)}$$

where $\mathbf{W}$ is a diagonal weighting matrix calibrated to emphasize primary driving variables.

Rather than deterministically selecting the closest match, the next state is sampled probabilistically from the k nearest neighbors using a non-parametric kernel probability distribution, such as the discrete Lall and Sharma kernel [57]. This probabilistic sampling introduces stochastic diversity while ensuring that synthesized values originate from physically observed multi-parameter states.

2. Kernel Density Estimation (KDE):

To prevent pure sample repetition and allow bounded exploration of unobserved states within the historical support, continuous non-parametric smoothing is applied via Kernel Density Estimation (KDE). KDE enables continuous resampling of non-Gaussian variables (such as skewed wind ramp rates and bimodal clearness indices) without imposing rigid parametric forms.

Several research papers employ KDE to generate synthetic meteorological and weather time series, primarily by using KDE as a non-parametric tool within stochastic weather generators to capture complex variable distributions without assuming a standard bell curve. Pavlides et al. [60] used a non-parametric, kernel-based estimator to predict precipitation amount. Gai et al. [61] used Gaussian KDEs to calculate the joint probability distribution

functions of dry/wet-bulb temperatures and solar radiation, leading to a synthetic "coincident design day" which accurately replicated actual extreme weather time series dynamics. Smith et al. [62] utilize Kernel Density Estimation and k-nearest-neighbor bootstrapping to generate synthetic time series of precipitation and temperature.

3. Block-Bootstrapping Strategies

To preserve multiscale temporal autocorrelation across meteorological variables without explicit parametric equations, block-bootstrapping methods resample consecutive contiguous chunks of historical observations to generate synthetic data. Standard implementations include the following:

- Moving Block Bootstrap (MBB): Resamples fixed-length blocks of length b with replacement from the historical time series [63–65].
- Stationary Bootstrap: Varies the block length L stochastically according to a geometric distribution $\mathbb{P}(L = l) = (1 - p)^{l-1}p$, ensuring that the resampled pseudo-time series remains strictly stationary [66].

Because entire multi-variable blocks are drawn from historic data, block-bootstrapping preserves internal cross-correlations across meteorological variables and intra-block temporal momentum across synoptic weather regimes.

4. Spectral methods

Within non-parametric surrogate modeling, ref. [67] developed a Discrete Fourier Transform (DFT) phase-manipulation framework (stsg) that generates synthetic environmental time series with controlled similarity while analytically preserving the empirical signal's first two statistical moments and autocorrelation structure.

2.3.2. Explicit Dependency Modeling via Copula Theory

While resampling preserves observed joint structures, it cannot decouple marginal behaviors from their underlying cross-variable dependence structures. Copula theory addresses this limitation by separating the modeling of individual variable distributions from their joint stochastic dependence.

The cornerstone is Sklar's Theorem [68], which allows splitting of multi-variable data into two independent parts: the individual shape of each variable (its margins) and the relationship connecting them (the copula). In synthetic data generation, this decouples how variables behave alone from how they move together. The process follows three steps: first, convert real observations into rank percentiles (0% to 100%) to strip away physical units; second, fit a copula function that models only the correlation structure between those percentiles; third, draw new random percentile pairs from the copula and map them back into real-world values using the individual distributions. This ensures synthetic data retains both individual profiles and realistic joint dependencies.

Standard copulas mainly fall into two families [69]: elliptical and Archimedean. Elliptical copulas, such as Gaussian and Student-t, model symmetric positive or negative linear correlations, with Student-t also capturing joint extreme spikes and dips. In contrast, Archimedean copulas (like Clayton, Gumbel, Frank), defined by a single generator equation, capture asymmetric relationships. For instance, Gumbel copulas link concurrent highs—such as blazing heat paired with intense sunshine—while Clayton copulas connect simultaneous lows, such as calm winds during cloudy spells.

Standard multivariate Archimedean and Elliptical copulas lack flexibility in higher dimensions ($d \geq 3$) because a single parameter or global correlation matrix is forced to govern all bivariate relationships. In weather synthesis involving $[GHI, v, T]$, different pairs exhibit distinct tail behaviors (e.g., wind and solar exhibit asymmetric joint distribution structures, whereas solar and temperature exhibit strong non-linear positive coupling).

Pair-Copula Constructions (PCCs), or Vine Copulas, resolve this dimensionality bottleneck by decomposing a d-dimensional copula density into a hierarchical product of $d(d-1)/2$ bivariate copulas (trees) conditioned on subsets of variables [70,71].

Using Canonical vines (C-vines) or Regular vines (R-vines), synthetic generators can assign distinct bivariate copula families to each pair in the tree (e.g., a Clayton copula for solar-wind drought pairs and a Gumbel copula for solar-temperature extremes). This formulation accurately captures joint non-linear cross-correlations and tail dependencies across multi-variable climate domains [72].

Alluri [73] used an R-vine copula to accurately capture complex dependence structure among wind power forecasts and reported surpassing the limitations of standard Gaussian and restricted vine copula models. Lazović and Đurišić [74] use a copula to capture the uncertainties of Day-Ahead electricity prices and wind park power generation, for the stochastic scheduling of a virtual power plant that integrates wind power and thermal storage for district heating.

Yang et al. [75] used a (Drawable/Decomposable) D-vine copula to address the challenges of the stochasticity of renewable generation on decentralized microgrid dispatch.

Jiao et al. [6] used Seasonal and Trend decomposition using Loess (STL) for the periodic and aperiodic components and then used an S-vine copula, before expanding the series and generating a 365-year synthetic year for Solar radiation, absolute humidity, and air temperature.

Sakki et al. [19] keep the classical treatment of deterministic periodicities and replace the remaining machinery with an explicit coupling of statistics, stochastics, and copulas. Claassen et al. [20] extend this to the spatial case using vine copulas, reproducing joint temperature, wind, and precipitation behavior at continental scale without assuming any parametric form for the dependence; though at daily resolution and without irradiance.

2.3.3. Strengths and Limitations of Non-Parametric Synthesis

Non-parametric synthesis offers distinct practical strengths along with a few notable limitations. On the plus side, it is extremely flexible because it adapts directly to the natural patterns of real-world observations. Data are not forced into rigid bell curves, which makes it easier to capture odd patterns like zero sunshine at night or complex, shifting wind conditions. Because the method draws straight from actual historical records or relative rankings, generated values always stay within real-world physical boundaries; meaning sunlight cannot go below zero, and relative humidity will not exceed 100%. In addition, techniques like nearest-neighbor lookups and copulas rely on straightforward geometric distances and mathematical optimizations. This keeps computation fast and mathematically transparent, avoiding the slow training runs, expensive graphics cards, and complex tuning associated with deep neural networks. On the downside, basic resampling cannot generate values higher or lower than what has already been recorded in history unless paired with specialized extreme-value math. Furthermore, stitching together sampled multi-day blocks often creates artificial jump discontinuities at the seams. These sudden jumps break the natural momentum of weather, like causing instantaneous temperature leaps or unrealistic wind gusts. As a result, non-parametric generators need careful boundary smoothing or interface matching to ensure smooth transitions from one day to the next.

2.4. Paradigm III: Deep Generative Architectures and Sequence Models

In recent years, the rapid maturation of deep learning has propelled neural network architectures to the forefront of synthetic time-series generation for power systems and climate impact studies [76,77]. Deep Generative Models (DGMs) represent a data-driven shift in synthetic time series generation. Rather than decomposing time series into analytical

trend-residual components or relying on explicit copula families, DGMs leverage deep neural network architectures to implicitly approximate high-dimensional joint probability distributions $p(\mathbf{X}_{1:T})$ from empirical data without structural assumptions.

Recurrent neural networks (RNNs), Long Short-Term Memory (LSTM) networks, Temporal Convolutional Networks (TCNs), and self-attention Transformer mechanisms are increasingly adopted by researchers and system operators seeking to capture non-linear, multi-scale temporal dependencies in high-resolution meteorological and renewable energy profiles. These sequence models have gained widespread traction precisely because they bypass manual feature engineering, automatically learning complex temporal hierarchies directly from massive reanalysis repositories and gridded sensor streams.

2.4.1. Generative Frameworks in Weather Synthesis

Neural networks map the inputs to the outputs through one or multiple layers of weighted, interconnected neurons [21,78–80]. There are several types of neural networks (e.g., Self-Organizing Map (SOM) [78,81], Convolutional Neural Network (CNN), Recurrent Neural Network (RNN), etc.). However, the ones with the most interest are Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs).

Generative Adversarial Networks (GANs) use two interconnected NNs (a Generator network G_θ and a Discriminator network D_ϕ) where the generator creates data using random noise and the discriminator applies a continuous value to distinguish between generated and real data [82]. In renewable and meteorological scenario generation, early implementations applied Deep Convolutional GANs (DCGANs) to generate 24 h wind and solar power trajectories [83]. To stabilize adversarial training and mitigate mode collapse, the Wasserstein GAN with Gradient Penalty (WGAN-GP) uses a different loss function [84,85]. Other notable GANs are Time-series GAN (TimeGAN) [86] and DoppelGANger [87] that attempt to capture complex temporal dynamics and have a wide range of applications. Hu et al. [88] used a conditional dual-attention generative adversarial framework (CTG-CSA) to improve modeling of multivariate meteorological dependencies across different climate regions.

GANs have been used to mitigate the scarcity of historical data for building energy studies. Tian et al. [89] used GANs to improve the accuracy of energy consumption prediction and to solve the problem of insufficient historical data, based on the parallel learning theory. Ma et al. [90] used an Attention-Based Cycle-Consistent Generative Adversarial Network (ABC-GAN) to generate IoT data for smart energy systems. Baasch et al. [91] used a conditional GAN to generate data suitable for retrofit analysis, smart grid integration, optimization, and load forecasting. Zhang et al. [92] combined GANs with wind turbine physical models to achieve both diversity and physical consistency in generated scenarios. Yang and Fu [93] reviewed recent trends in PV variability and compared explicit probabilistic/statistical models, deep generative models, and Statistical Relational AI (SRAI)–integrated frameworks.

Variational Autoencoders map the input meteorological series $\mathbf{x}$ using an encoder network in a latent space (thus learning its probability distribution), while a decoder network $p_\theta(\mathbf{x}|\mathbf{z})$ reconstructs it to the original space [94,95]. The model is optimized by maximizing the Evidence Lower Bound (ELBO). Fan et al. [96] introduced the use of VAEs for improving short-term building energy predictions. Bregere et al. [97] used a different subtype, Conditional Variational Autoencoders (CVAE), for the simulation of Tariff impact on electric energy consumption.

Denoising Diffusion Probabilistic Models (DDPMs) and score-based generative models are a class of likelihood-based models that have demonstrated remarkable results in computer vision [98,99]. They work by corrupting historical sequences through a forward

Markovian Gaussian noising process, and a neural network is then trained to iteratively reverse this perturbation using score matching. Capel et al. [100] used DDPMS for probabilistic Energy forecasting using the open data of the Global Energy Forecasting competitions, and found the results comparable to state-of-the-art deep learning generative models, including generative adversarial networks, variational autoencoders, and normalizing flows.

Frameworks such as TimeGrad [101] and Conditional Score-based Diffusion Imputation (CSDI) [102] combine diffusion reversing kernels with Recurrent Neural Networks (RNNs) and bi-directional self-attention transformers [101,102]. These architectures capture non-local multi-scale temporal dependencies, establishing state-of-the-art benchmarks for complex conditional scenario synthesis.

2.4.2. The "Short-Window" Bottleneck in Climatological Generation

Despite their strong representational power, deep generative models face structural bottlenecks when scaled from short-horizon benchmarks to continuous annual climatological simulations (L = 8760 h), as illustrated in Figure 1.

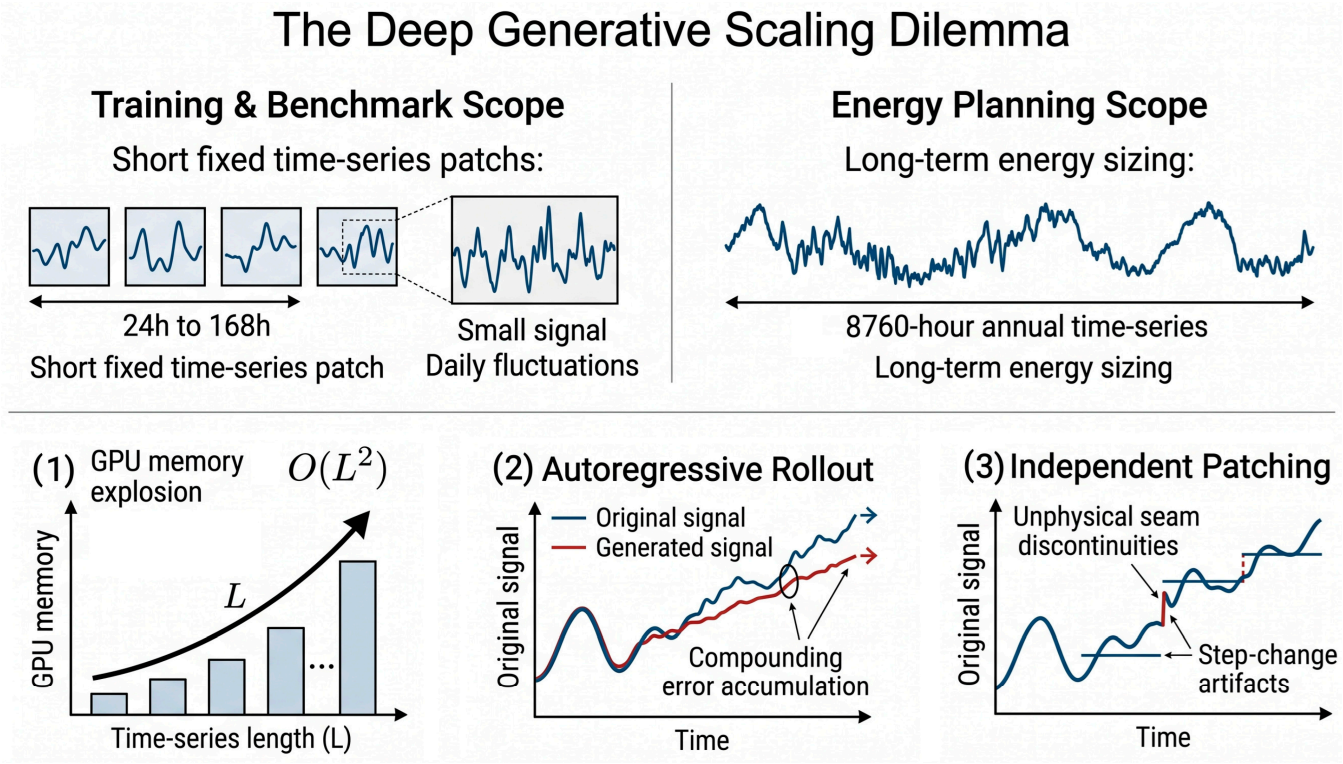


Figure 1. The deep generative scaling dilemma.

A primary constraint stems from hardware capacity and computational scaling. Standard self-attention mechanisms scale quadratically with sequence length ($O(L^2)$). While tractable over daily (L = 24) or weekly (L = 168) intervals, extending training directly across an entire continuous hourly annual sequence (L = 8760) rapidly exhausts GPU memory. This scaling barrier necessitates massive parameter spaces that are prone to overfitting sparse multi-decadal historical catalogs.

Generating annual records through rolling autoregressive sampling introduces compounding errors. Under open-loop execution, minor pointwise discrepancies accumulate over successive simulation steps. Over thousands of consecutive hours, this error propagation induces severe trajectory drift, causing synthetic series to diverge into unrealistic variances or collapse toward static attractors.

A common workaround—independently generating short daily or weekly segments and concatenating them into an annual series—introduces seamline artifacts (Patch-Stitching Seamline Discontinuities). Unconditioned boundary transitions disrupt temporal continuity, creating unphysical step discontinuities across patch interfaces (e.g., t = 24,48,72,...). These boundary jumps distort high-frequency ramp rates, corrupt power spectral distributions, and skew simulated storage charging dynamics.

2.4.3. Physical Inconsistencies and Computational Overheads

Unlike parametric and non-parametric empirical methods, standard deep generative models lack inductive physical biases and may violate physical bounds or conservation laws. Because neural activation layers (e.g., linear, ReLU, sigmoid) operate purely on unconstrained statistical approximations, unconstrained DGMs frequently generate the following:

- Nocturnal Solar Inconsistencies: Generating non-zero positive global horizontal irradiance ($GHI > 0$) during nighttime hours.
- Thermodynamic Support Violations: Synthesizing relative humidity values outside physical bounds ($RH < 0\%$ or $RH > 100\%$) or dry-bulb temperatures that drop below corresponding dew-point temperatures.
- Unnatural Wind Dynamics: Simulating negative wind velocities or instantaneous acceleration spikes that violate aerodynamic boundary-layer momentum equations.

Enforcing these bounds through ad hoc post-processing clipping (e.g., $\max(0,GHI)$) creates artificial probability spikes at boundaries, corrupting the underlying generated probability density.

Deep generative models introduce significant computational overhead and usability challenges in practical workflows [103]. Generative adversarial networks often face training instability, including vanishing gradients and non-convergent minimax cycles that demand exhaustive hyperparameter optimization across update ratios, regularizers, and learning rates [103,104]. Conversely, diffusion models exhibit substantial sampling latency, as their iterative reverse denoising process requires evaluating tens to hundreds of numerical steps per sample [102]. This creates a severe computational bottleneck when simulating large Monte Carlo ensembles over extended annual horizons. Compounding these resource demands is the black-box nature of deep architectures, whose vast, high-dimensional parameter spaces lack analytical transparency [105]. Consequently, identifying and correcting the mathematical causes of anomalous or unphysical synthetic scenarios proves challenging, particularly when compared to the interpretable formulation of distribution-free non-parametric frameworks.

3. Mathematical Methodology and Modeling Framework

3.1. Overview of MeteoSynth Pipeline

MeteoSynth generates synthetic hourly meteorological series by combining a time-inhomogeneous first-order Markov chain with non-parametric, kernel-based transition densities. An overview of the MeteoSynth pipeline is presented in Figure 2. For each target calendar day, training data are pooled from a rolling window of $\pm n$ days (default $n = 15$) centered on that day (Figure 2a). The window is evaluated on a circular calendar, so that it wraps across the New Year, and draws from all years of the historical record. The pooled days are then arranged as a days × hours table, and adjacent hourly columns are paired to form 23 within-day transition samples (x_{h-1}, x_h) where $h = 1, \ldots, 23$ (Figure 2b). Together with the unconditional initial distribution $\pi(x_0)$ at 00:00 and a cross-midnight transition from 23:00 to 00:00 of the following day, this yields 25 models per calendar day (although the

unconditional initial distribution is only required for the very first day of a multiday cycle). Each transition sample is fitted with a bivariate Gaussian kernel density estimate of the joint density $\hat{p}(x_{t-1}, x_t)$ (Figure 2c). Given the current state $x_{t-1} = v$, the joint density is sliced at v and renormalised, yielding the conditional density $p(x_t|v)$ and its cumulative distribution $F(x_t|v)$ (Figure 2d), from which the next state is drawn by inverse-transform sampling, $x_t = F^{-1}(\mathrm{u} \mid \mathrm{v})$ with $\mathrm{u} \sim \mathcal{U}(0, 1)$ (Figure 2e). The drawn value becomes the conditioning value for the following hour; consecutive days are linked through the cross-midnight model, and the training window advances one day at a time, producing continuous multi-day synthetic series (Figure 2f).

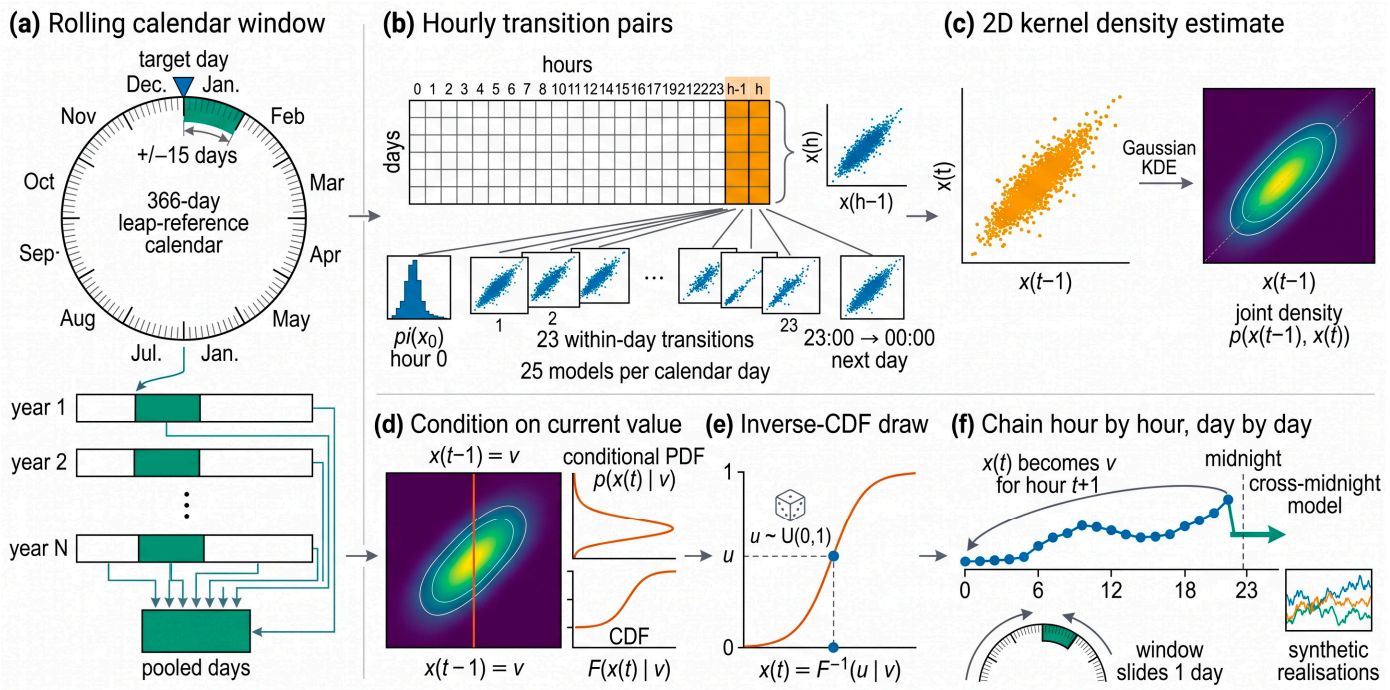


Figure 2. Schematic of the MeteoSynth generation pipeline.

Section 3.2 describes this core generative model in the order in which data flow through it. Section 3.3 presents the safeguards that make it physically consistent and numerically robust on real weather records, and Section 3.4 the implementation choices that make large Monte-Carlo ensembles tractable.

3.2. Core Generative Model

3.2.1. Climatological Training Window

Markov chain training requires selecting a representative subset of historical data. The subsetting is presented in Figure 2a. MeteoSynth implements two distinct subsetting strategies: Monthly Subsetting and Rolling Day Window (default).

The Monthly Subsetting strategy pools all days belonging to a specific calendar month (e.g., 1–31 May). While computationally much more efficient, monthly subsetting exhibits two major drawbacks: (A) Off-center bias: A model fitted for 1 May trains on data centered on 16 May, biasing early-month simulations. (B) Boundary step discontinuities: Disjoint training pools between 31 May and 1 June cause unphysical statistical jumps across month boundaries.

The Rolling Day Window is the default strategy. Pools all days falling within $\pm n$ days of the target calendar date across all recorded years. For example, a window half-width of

$n = 15$ days over an 11-year record pools $(2 \times 15 + 1) \times 11 = 341$ days, ensuring centered, smooth seasonal transitions.

In order to handle leap years and New Year wrapping without pandas timestamp overhead, MeteoSynth implements a 366-day circular reference calendar (the all_leap calendar of the Climate and Forecast (CF) metadata conventions [106], the standard for describing climate data, in which every year contains 29 February). Each calendar date (month, day) is mapped to an integer slot $s \in [1,366]$ (where 29 February is permanently fixed to slot 60, and 1 March is slot 61, etc.). Distance between target slot c and sample slot s is evaluated using a circular distance metric with period 366:

$$\delta(s,c) = \min((s-c) \pmod{366},\ 366 - ((s-c) \bmod 366))$$

An observation day is included in the fitting window if and only if $\delta(s,c) \leq n$. This formulation provides two key operational features. First, it ensures New Year wrapping: a target slot of 2 January ($c = 2$) with a window half-width of $n = 15$ naturally incorporates historical observations from late December ($s = 352 \ldots 366$) without requiring manual year-boundary conditional logic. Second, it provides robustness across leap years: non-leap years simply omit slot 60 without disrupting the index alignment of March through December.

The sampling rules for 29 February follow directly from this construction, and none of them duplicates or reweights an observation:

- Training. Every observed day enters a window once, at its own slot. An observed 29 February exists only in leap years (4 of the 19 years of the case study of Section 4), so it is never repeated to fill the other years. A window that covers slot 60 is simply 15 days smaller: with n = 7 it pools $14 \times 19 + 4 = 270$ days instead of 285. Each 29 February observation thus carries the same weight as any other day, and its share of a pool is about 1.5%.
- Cross-midnight pairs: These are built from real consecutive dates (Section 3.2.5): 28 February → 29 February and 29 February → 1 March in leap years, 28 February → 1 March otherwise.
- Generation. By default, a synthetic year has 365 days, and the calendar walk skips slot 60, stepping from 28 February 23:00 to 1 March 00:00 with the cross-midnight model of 1 March, whose pairs include both kinds of real predecessors. Setting include_feb29 = True produces 366-day years, in which 29 February is generated from its own window, centered on slot 60 and dominated by the neighboring days of late February and early March. If a record contains no leap year at all, that window is formed from the neighboring days alone (for any n ≥ 1), so a synthetic 29 February can still be generated without runtime errors.

3.2.2. Stochastic Formulation and First-Order Markov Assumption

Figure 2b presents the next step in the process. Let $X_t \in \mathbb{R}$ represent a continuous meteorological variable (e.g., global solar irradiance G, dry-bulb air temperature T_{air}, or wind speed v_{wind}) observed at hourly discrete time steps $t \in \{0,1,\ldots,23\}$ across consecutive days. The evolution of X_t during the day is modeled as a first-order, time-inhomogeneous Markov process. The conditional probability distribution of the state at step t depends strictly upon the state at the immediately preceding step $t-1$:

$$P(X_t = x_t \mid X_{t-1} = x_{t-1}, X_{t-2} = x_{t-2}, \ldots, X_0 = x_0) = P(X_t = x_t \mid X_{t-1} = x_{t-1})$$

Because diurnal solar trajectory and thermal inertia dictate different statistical behaviors at different hours of the day, a single stationary transition kernel cannot capture diurnal microclimates. Consequently, MeteoSynth constructs time-inhomogeneous transition models: for every calendar day, 23 distinct transition kernels are fitted independently for each hourly step $h \in \{1,2,\ldots,23\}$, mapping hour $(h-1) \rightarrow h$.

The training pairs $\left\{\left(x_{h-1}^{(i)}, x_h^{(i)}\right)\right\}_{i=1}^{n}$ for each kernel are read from adjacent columns of the days × hours table built from the training window of Section 3.2.1, which guarantees that both members of a pair originate from the same day. Hour 0 has no within-day predecessor and is handled separately (covered in Section 3.2.5).

3.2.3. Non-Parametric 2D Kernel Density Estimation (2D KDE)

Figure 2c shows how each partitioned transition matrix is transformed through the KDE 2d Estimation. To model the joint probability density function $\hat{p}(x_{t-1}, x_t)$ without imposing restrictive parametric shapes, MeteoSynth applies a two-dimensional Gaussian Kernel Density Estimator to historical observation pairs $\left\{\left(x_{t-1}^{(i)}, x_t^{(i)}\right)\right\}_{i=1}^{n}$ extracted from a pooled training window:

$$\hat{p}(x_{t-1}, x_t) = \frac{1}{n|\boldsymbol{H}|^{1/2}} \sum_{i=1}^{n} K_{\boldsymbol{H}}\left(\begin{bmatrix} x_{t-1} \\ x_t \end{bmatrix} - \begin{bmatrix} x_{t-1}^{(i)} \\ x_t^{(i)} \end{bmatrix}\right)$$

where $K_{\boldsymbol{H}}(\boldsymbol{z}) = \frac{1}{2\pi} \exp\left(-\frac{1}{2}\boldsymbol{z}^T \boldsymbol{H}^{-1} \boldsymbol{z}\right)$ is a standard bivariate Gaussian kernel, and $\boldsymbol{H}$ is the 2×2 bandwidth covariance matrix. Bandwidth matrix selection governs the trade-off between over-smoothing and overfitting. MeteoSynth supports automated bandwidth estimation via Scott's rule. However, empirical diagnostics revealed that Scott's rule systematically over-smooths solar irradiance transitions and therefore a calibrated fixed relative bandwidth factor ($h = 0.1$) was used, which preserves sharp transitions while avoiding sample overfitting; however, a user has the option to adjust that value.

3.2.4. Conditional Sampling by Inverse Transform

To sample the next state x_t given a known current state $X_{t-1} = v$, MeteoSynth takes a continuous slice through the joint density function at $x_{t-1} = v$ and normalizes it to obtain a proper one-dimensional conditional probability density function:

$$p(x_t \mid X_{t-1} = v) = \frac{\hat{p}(v, x_t)}{\int_{S_X} \hat{p}(v, x)\, dx}$$

where $S_X \subset \mathbb{R}$ represents the valid physical support of variable X. The slicing process is presented in Figure 2d.

The next step is the sampling of the new state value, which is shown in Figure 2e. Sampling is performed via inverse-transform sampling; i.e., the cumulative distribution function (CDF) $F(x_t \mid v) = \int_{-\infty}^{x_t} p(x \mid v) dx$ is constructed over a discrete evaluation grid; a uniform random variable $u \sim \mathcal{U}(0, 1)$ is drawn, and the synthetic value for time t is evaluated via numerical quantile inversion:

$$x_t = F^{-1}(u \mid v)$$

3.2.5. Chaining Across Hours and Days: π(X_0) and Cross-Midnight Continuity

Finally, the chaining between days is presented in Figure 2f. Within a day, the value drawn at hour t becomes the conditioning value v for the hour $(t + 1)$ transition model, so that a complete diurnal profile is produced by walking the 23 hourly kernels in sequence. Because hour 0 (00:00) has no preceding hour within the same calendar day, a multi-day Markov chain requires specialized boundary handling:

1. Initial Distribution $\pi(X_0)$: When starting a fresh synthetic series, hour 0 is sampled unconditionally from $\pi(X_0)$, an empirical 1D distribution fitted to all observed 00:00 values within the target date's training window. For solar irradiance, $\pi(X_0) \equiv 0$. For air temperature and wind speed, $\pi(X_0)$ captures realistic midnight state distributions. This is only implemented once for each realization.

2. Cross-Midnight Model ($P(X_{0,\text{today}} | X_{23,\text{yesterday}})$): When chaining multi-day synthetic time series, setting hour 0 of day d from $\pi(X_0)$ would destroy inter-day serial correlation. Instead, MeteoSynth fits a dedicated cross-midnight model mapping X_{23} of day $d-1$ to X_0 of day d. It is important that the transition pairs are extracted strictly from consecutive calendar dates in the historical record rather than arbitrary window slots, preventing spurious multi-year temporal jumps.

Each calendar day is therefore described by 25 models: $\pi(X_0)$, the 23 within-day transitions, and the cross-midnight transition. As the series advances from day d to day $d+1$, the training window of Section 3.2.1 slides by one slot, so the chain is time-inhomogeneous at daily as well as hourly resolution and the fitted models drift smoothly through the season.

3.3. Physical and Numerical Safeguards

3.3.1. Physical Support Truncation

A common flaw in standard KDE implementations is boundary leakage: Gaussian kernels have infinite tails and extend past natural boundaries, thus creating non-physical synthetic values, such as negative solar irradiance ($G < 0$) or negative wind speeds ($v_{\text{wind}} < 0$). Common ad hoc fixes (such as post hoc zero-clipping or 10% domain padding) distort the integral of the distribution, leading to artificial probability spikes at zero or out-of-domain sampling.

MeteoSynth separates the two pieces of information that together define where a value may lie. The data define where the mass is: the observed range $[x_{\text{obs,min}}, x_{\text{obs,max}}]$ of the training sample (Figure 3a). Physics defines where a value is allowed to be at all: each variable declares its support $S_X = [x_{\min}, x_{\max}]$, together with its units and long name, once, in a VariableSpec registered in meteosynth.variables and presented in Table 1 and represented in Figure 3b. The supported range is a hard limit that the data cannot override; the data decide how much of the support is actually used.

Table 1. Meteorological variables—units and physical boundaries.

Variable Name	Attribute Key	Physical Support (S_X)	Units	Description
Solar Irradiance (POA)	poa_global	$[0, +\infty)$	$\frac{\text{W}}{\text{m}^2}$	Plane-of-Array Irradiance
Solar Irradiance (Horizontal)	ghi	$[0, +\infty)$	$\frac{\text{W}}{\text{m}^2}$	Global Horizontal Irradiance
Air Temperature	temp_air	$[-273.15, +\infty)$	°C	Dry-Bulb Air Temperature
Wind Speed	wind_speed	$[0, +\infty)$	m/s	Wind Speed at 10 m Height

The VariableSpec registry is also what makes MeteoSynth variable-agnostic. A column can be modeled if and only if it has a registered specification. Table 1 lists the specifications that ship with the package, and a user can add a new one with a single call. For example, registering relative humidity with support [0,100] and units of % is enough, with no other change to the package. The upper bound is then enforced in the same way as the lower bounds discussed here. The three variables studied in this paper were chosen to cover the main kinds of support: a lower bound that is effectively never reached (temperature), a lower bound that is approached but rarely touched (wind speed), and a lower bound that the data sit on for long periods (irradiance, which is also degenerate at night; Section 3.3.3).

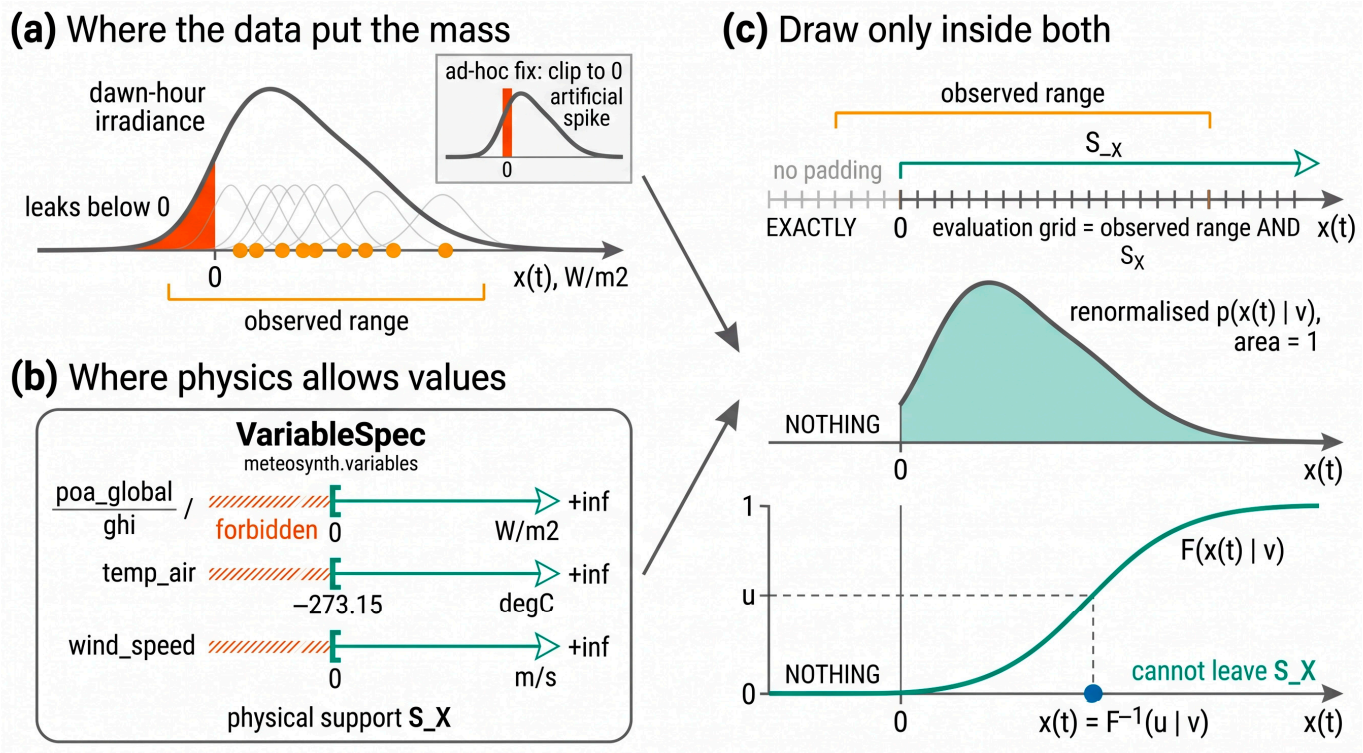


Figure 3. Physical support truncation. (**a**) A Gaussian KDE fitted to dawn irradiance leaks mass below zero; clipping the draws instead piles it into a spike at the bound. (**b**) Each variable declares its physical support S_X once. (**c**) The conditional density is evaluated only on the intersection of the observed range and S_X, so inverse-CDF draws cannot leave the support.

The conditional density of Figure 3c is then evaluated only on the intersection of the two:

$$[x_{\text{lo}}, x_{\text{hi}}] = [x_{\text{obs,min}}, x_{\text{obs,max}}] \cap S_X,$$

discretized as a grid of 1000 equally spaced points with no padding on either side. The slice $\hat{p}(v, x_t)$ is evaluated at these points and renormalized so that its area on $[x_{\text{lo}}, x_{\text{hi}}]$ is one; its cumulative sum gives the discrete CDF $F(x_t|v)$, which rises from 0 at x_{lo} to 1 at x_{hi}. The drawing of a new value $x_t = F^{-1}(u|v)$ is a linear interpolation of the grid against these CDF values. Because linear interpolation can only return values between the grid points it is given, and the grid lies inside S_X by construction, an out-of-support sample is impossible. No clipping of the output is needed and no spike is created: the mass that would have leaked below the bound is simply redistributed over the admissible range by the renormalisation.

3.3.2. Clamping the Conditioning Value at Model Hand-Over

A separate numerical guard acts on the *input* of a transition model, rather than on its output. Slicing the joint density at $x_{t-1} = v$ (Figure 4a) is only valid where the next state value v lies within the range of x_{t-1} values that the kernels were fitted on. If v is outside of the range, the integral vanishes, and the conditional density becomes an indeterminate $0/0$ that would manifest as a non-finite value that would contaminate the rest of the series. Such a value v that exists outside the range can potentially be obtained at any hand-over from one hourly model to the next; however, it is most likely at the transition between one day and the next (Figure 4a).

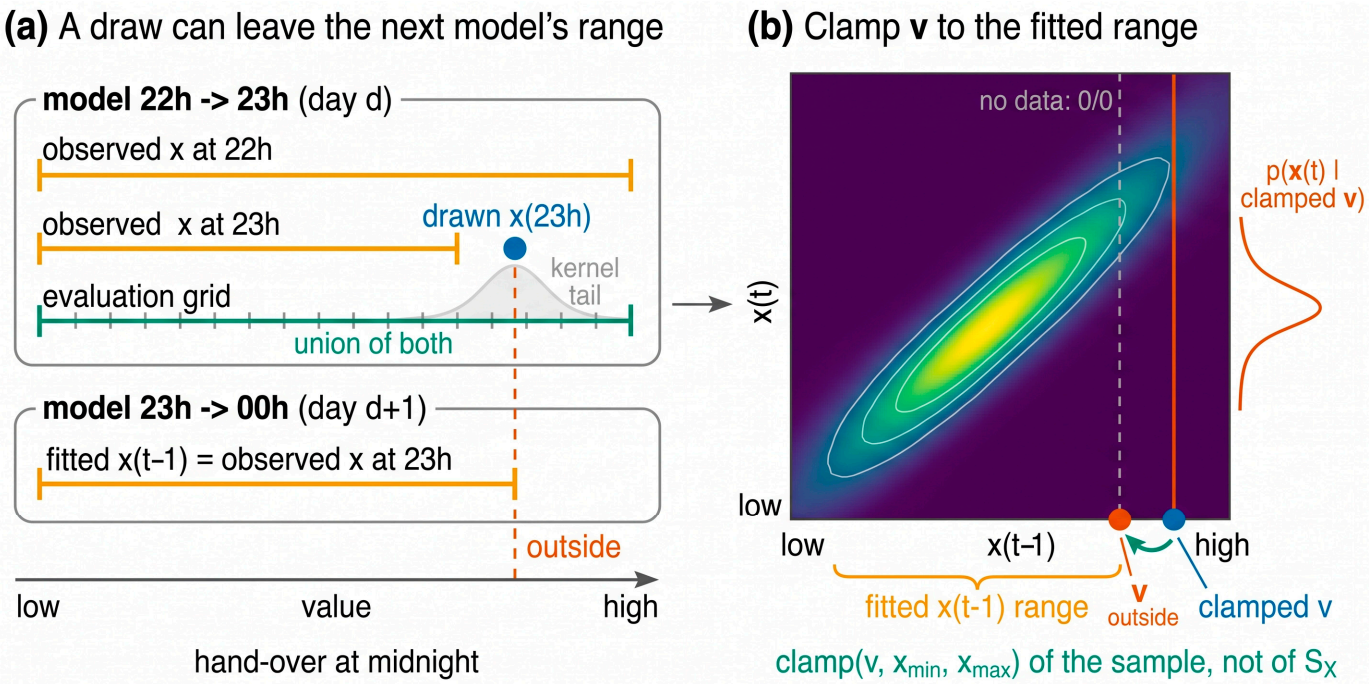


Figure 4. Clamping the conditioning value at model hand-over. (**a**) The 22 h → 23 h model draws on a grid spanning both hours' ranges, so a 23:00 draw can exceed every observed 23:00 value; the next day's 23 h → 00 h model, might have no data there. (**b**) The value is clamped to the fitted x(t − 1) range before slicing.

The evaluation grid of Section 3.3.1 is taken over both columns of the transition sample, so the 22:00 → 23:00 model draws on the union of the 22:00 and 23:00 observed ranges, and the kernel tails carry a little mass beyond the most extreme observed 23:00 value. The 23:00 → 00:00 model of the following day is fitted on (x_{23}, x_0) pairs whose $x_{\{23\}}$ column is exactly the observed 23:00 values, so a draw in that tail lies beyond its range. The samples of consecutive fitted models can also differ marginally at the record edges and around 29 February.

MeteoSynth therefore clamps the conditioning value to the nearest edge of the fitted sample range before slicing (Figure 4b), $v \leftarrow \mathrm{clamp}(v, x_{t-1,\min}, x_{t-1,\max})$. Since a KDE carries no information beyond its data, the nearest fitted value is both the safest and the most faithful conditioning point. The clamp is to the sample range, not to S_X; the two coincide only where the data touch the physical bound, as irradiance does at dawn. The event is rare: the excess mass lies in the outermost kernel tail beyond the most extreme observation, which scales inversely with the number of training pairs $(2n + 1)\,Y$ for a window half-width n and Y years of record, so it becomes rarer as the window widens.

Even though this scenario represents an exceedingly rare edge case, the guard is implemented because an unhandled occurrence—however improbable—would prematurely terminate the generation of the synthetic time series.

3.3.3. Degenerate State Classification (Kind)

Environmental time series frequently exhibit zero-variance regimes, such as solar irradiance during nighttime hours. Directly fitting a two-dimensional Gaussian kernel density estimator to transition pairs where both successive observations are invariant (where $x_{t-1} = 0$ and $x_t = 0$ causes the sample covariance matrix to become singular ($\det(\mathbf{\Sigma}) = 0$), resulting in numerical instability. Rather than checking for exact floating-point zeros—which is fragile when ingesting records with residual numerical noise from unit conversions—the

dispersion at each step is evaluated relative to the maximum axis scale $M = \max(|x_{t-1}|,|x_t|)$, using a relative tolerance $\epsilon = 10^{-6}$.

To prevent numerical singularities while accommodating noisy historical records, transitions are classified into three operational kinds based on the dispersion of the conditioning and target states relative to the maximum axis scale:

$$\text{Kind} = \begin{cases} \text{CONSTANT}, & \text{if} \sigma(x_t) \le \epsilon M, \\ \text{KDE_1D}, & \text{if} \sigma(x_t) > \epsilon M \text{ and } \sigma(x_{t-1}) \le \epsilon M, \\ \text{KDE_2D}, & \text{if} \sigma(x_t) > \epsilon M \text{ and } \sigma(x_{t-1}) > \epsilon M. \end{cases}$$

Under this classification (which can be seen in Figure 5), invariant target periods ($\sigma(x_t) \le \epsilon M$) bypass stochastic simulation entirely and deterministically assign the constant mean $\mu(x_t)$, effectively reproducing nocturnal zero solar irradiance periods. When the target state exhibits active dispersion but the previous state remains invariant ($\sigma(x_t) > \epsilon M$ and $\sigma(x_{t-1}) \le \epsilon M$), the system samples unconditionally from a one-dimensional density estimate fitted solely over x_t, as occurs during the night-to-dawn (early morning) transition. Finally, when both successive states exhibit sufficient variance ($\sigma(x_t) > \epsilon M$ and $\sigma(x_{t-1}) > \epsilon M$), the full two-dimensional conditional kernel density estimator (KDE_2D) is deployed. Should near-collinear observations induce ill-conditioning in the sample covariance matrix during inversion, the solver gracefully resolves to the unconditional one-dimensional estimate rather than terminating execution.

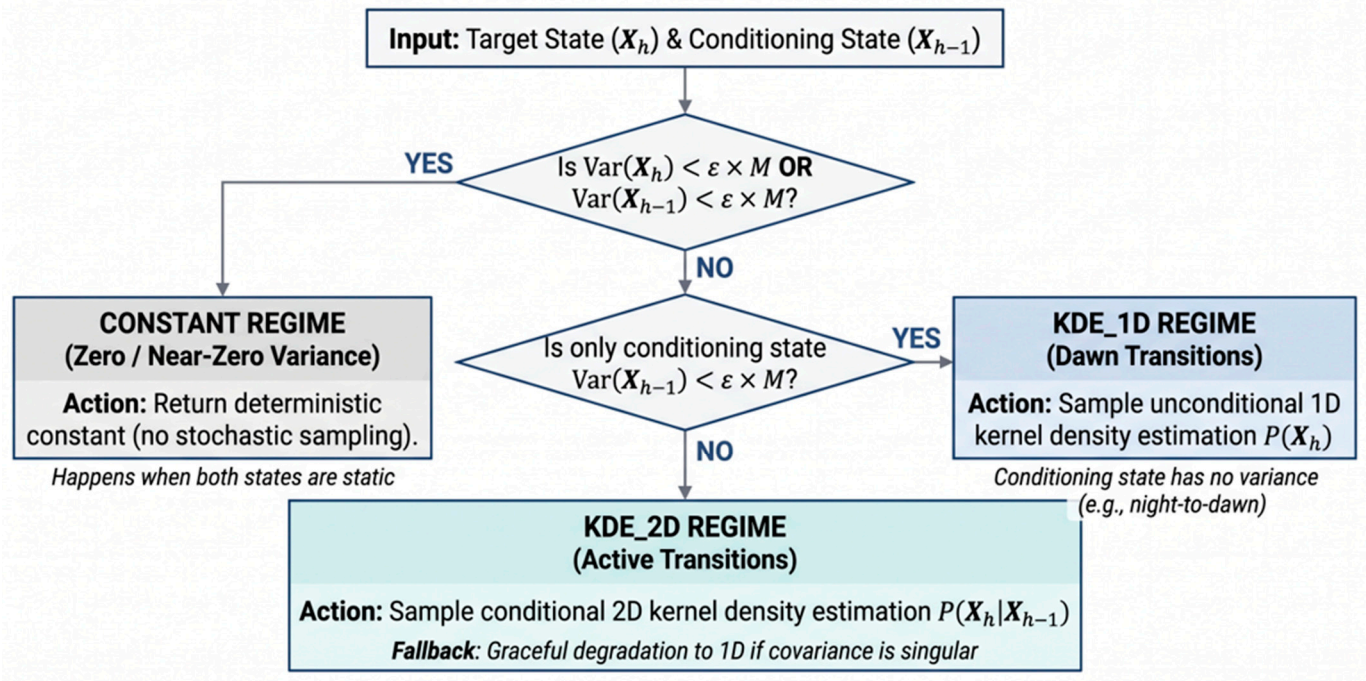


Figure 5. Decision logic for degenerate state classification and numerical stability in Markov transition estimation. Transitions between consecutive hourly states (x_{t-1},x_t) are screened against a relative variance threshold (ϵM) to prevent covariance matrix singularity ($\det(\mathbf{\Sigma}) = 0$), routing execution into deterministic constant return (CONSTANT), unconditional 1D sampling (KDE_1D), or full conditional 2D estimation (KDE_2D) with fallback degradation.

3.4. Computational Efficiency Refinements

3.4.1. Cached Conditional Quantile Grid

Evaluating a 1000-point 2D KDE slice for every hourly draw accounts for ~97% of total execution time. To overcome this computational bottleneck during large-scale Monte Carlo runs, MeteoSynth implements a lazy precomputed conditional quantile grid. When a

generator instance receives more than 64 draw requests or is called within ensemble mode, it constructs a two-dimensional lookup table of conditional quantile functions across a 64-node grid of conditioning states X_{t-1}. Subsequent sampling evaluates target values $x_t = F^{-1}(u|v)$ through bilinear interpolation: first by interpolating conditional quantile vectors along X_{t-1} across bracketing grid nodes and subsequently interpolating along the quantile probability axis $u \sim \mathcal{U}(0,1)$.

Performing interpolation directly on quantile functions $F^{-1}(u)$ rather than probability densities preserves unimodal distribution across varying conditioning locations. In benchmark evaluations, this cached quantile sampler achieves a thirteenfold speedup per sampled draw while preserving statistical fidelity, matching exact density evaluation within 0.03% on the sample mean and 0.31% on the standard deviation.

3.4.2. Single Calendar Walk for Ensemble Generation

The main computational cost of generating a next value is the step of the transition matrix and the fitting to the bivariate Kde2D. When creating an ensemble of 1000 full-year realizations, there are significant speed benefits if the fitting is performed only once, and the drawing occurs for all realizations, as shown in Figure 6. The benefits are greater as the number of realizations increases. For an ensemble of N = 1000 full-year realizations of one variable, single-walk generation is approximately two orders of magnitude faster than naive serial generation. On the machine tested, the run time fell from about 27 h to 21 min, a factor of 77. This figure is indicative only, since it depends on the hardware, the number of realizations, and the window size.

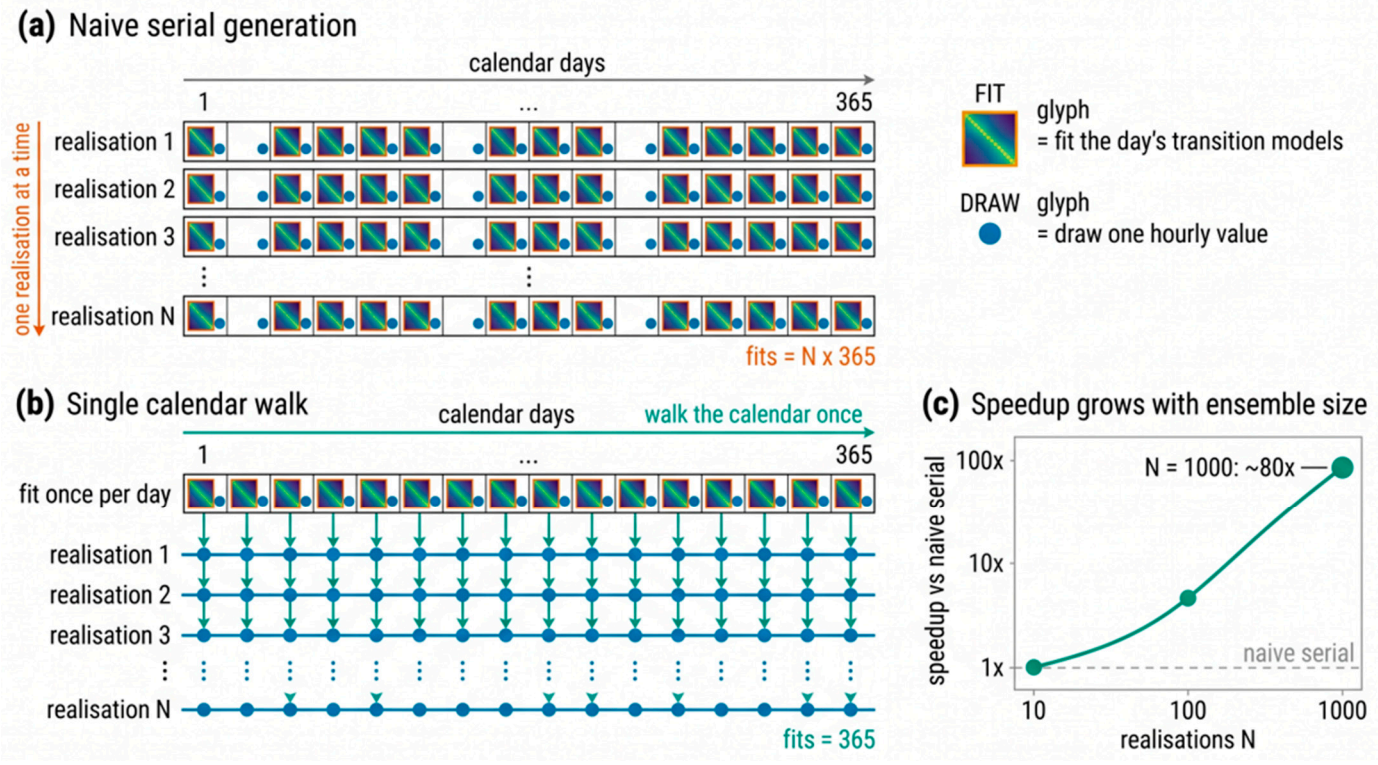


Figure 6. Ensemble generation optimization.

4. Results and Discussion

4.1. Methodology

The generator was tested for Heraklion, Crete (Lat: 35.3387, Long: 25.1442), using 19 years (2005–2023) of hourly PVGIS data for plane-of-array irradiance, air temperature, and

wind speed. The site was chosen for convenience, as the location of the author's department. It was not selected as representative of other climates, nor as particularly favorable or unfavorable for the generator. The PVGIS default plane was used, i.e., a horizontal surface (tilt 0°), so the plane-of-array irradiance analyzed here is the global horizontal irradiance including horizon shading (mean annual total 1879 kWh/m^2). Each variable is generated separately. Because all three follow the same calendar and hour of day, their shared daily and seasonal pattern is maintained: a July midday is still both sunny and hot. However, the link between weather departures from the norm is lost: in the observed record, a cloudy winter day also tends to be cooler and windier than usual, whereas in the synthetic data it does not. The observed correlations are reported in Section 4.4.5.

The evaluation is structured into two main experimental setups. First, Section 4.2 examines diurnal performance on a representative calendar day (21 May) using 2000 realizations generated with a ±15-day training window. Section 4.3 through Section 4.5 subsequently assess full annual dynamics using two ensembles of 1000 synthetic years each, configured with ±7-day and ±15-day training windows, respectively. The window parameter determines the temporal range pooled across all 19 observational years to calibrate each calendar day's transition model, yielding effective sample pools of 285 days for the ±7-day window and 589 days for the ±15-day window. While a wider window increases the sample size available for density estimation, it inevitably incorporates days further away in the season.

The annual validation framework evaluates model performance along two primary axes. First, marginal fidelity is assessed using the first-order Wasserstein distance (W_1) to determine whether synthetic distributions preserve the observed range and density across each month-by-hour cell. Second, temporal fidelity is evaluated using the autocorrelation function (ACF) to determine whether synthetic trajectories capture the temporal persistence of empirical weather patterns. For both evaluations, model fidelity is benchmarked directly against the natural inter-annual variability of the historical record, recognizing that real meteorological conditions naturally exhibit substantial dispersion across the 19-year baseline.

4.2. Typical Day Example

Before looking at whole years, the results for a single day are investigated and presented. A typical-day comparison shows directly whether synthetic days have the right shape, the right level, and the right day-to-day variety, and it is the kind of check a user can repeat for any date of interest.

Figure 7 compares synthetic and observed irradiance for 21 May. In the left panel, the thin blue lines are 21 May of each observed year, and the light-blue band is the 10th–90th percentile (P10–P90) of all 285 observed days within ±7 days of that date: at each hour, eight out of ten observed days fall inside the band. The thin orange lines are 6 of the 2000 synthetic days, and the dotted orange lines are their P10–P90.

The observed and synthetic medians (thick lines) almost coincide, and so do the upper P90 lines. Like real days, some synthetic days are smooth clear-sky curves, and others are broken by clouds. Around midday, the synthetic P10 lies somewhat above the observed band, so very dark, overcast hours are a little rarer in the synthetic data. The right panel shows the daily total: both distributions have a clear-sky peak near 8 kWh/m^2 and a long tail towards cloudy days, in close agreement.

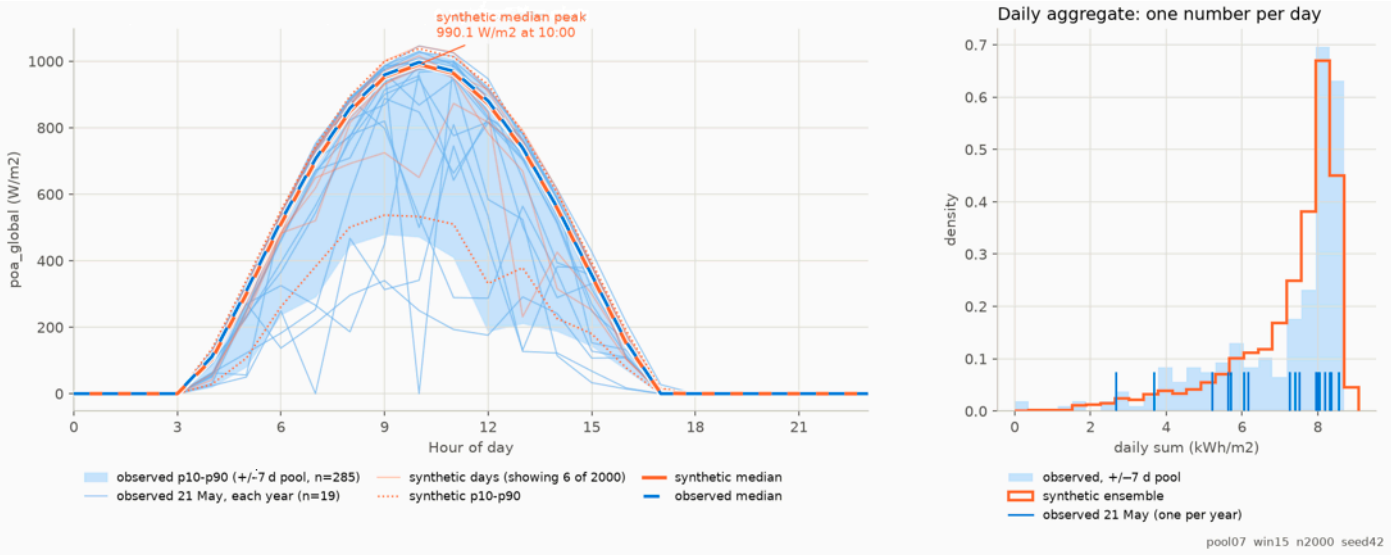


Figure 7. Typical day (21 May) at Heraklion, plane-of-array irradiance. (**Left**): hourly values, observed (blue) and synthetic (orange). (**Right**): distribution of the daily total.

For wind speed (Figure 8), the envelopes agree well, and the medians coincide, peaking at about 3 m/s around noon. Individual days, real and synthetic, vary much more from hour to hour than irradiance, because wind follows the day–night cycle less strictly. The daily mean is skewed to the right, with a long tail of windy days, and again the two distributions agree well.

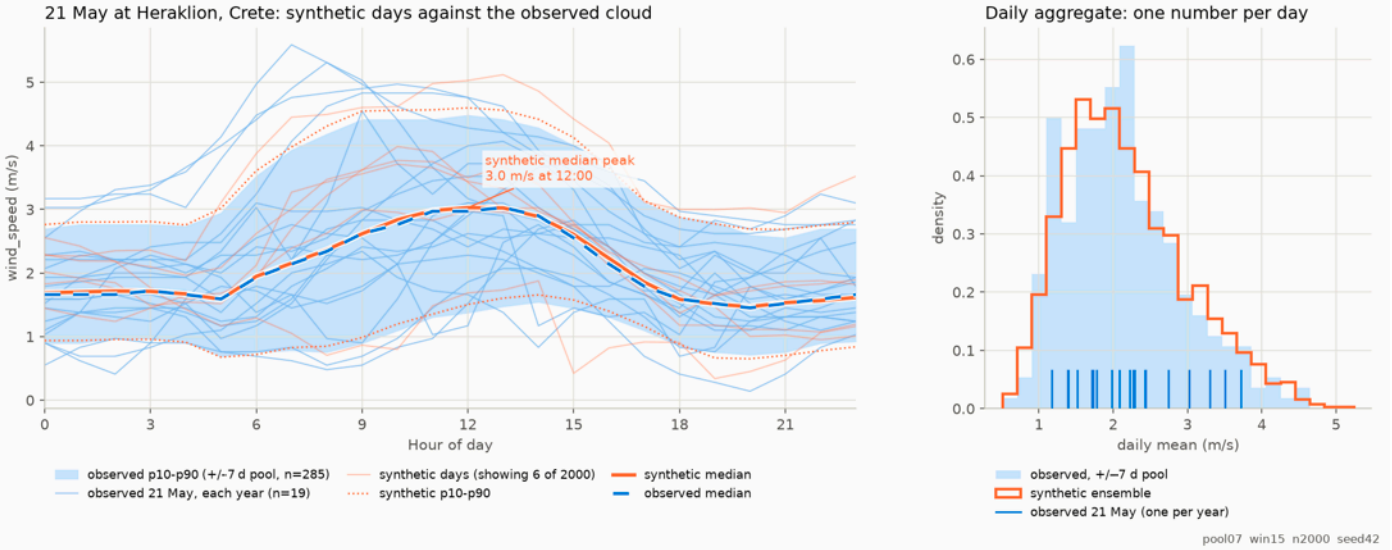


Figure 8. Typical day (21 May) at Heraklion, wind speed. **Left**: hourly values. **Right**: distribution of the daily mean.

For air temperature (Figure 9), the medians lie on top of each other throughout the whole day. The synthetic envelope follows the observed band but is slightly wider, mainly on the cool side, and the daily mean distribution shows the same picture.

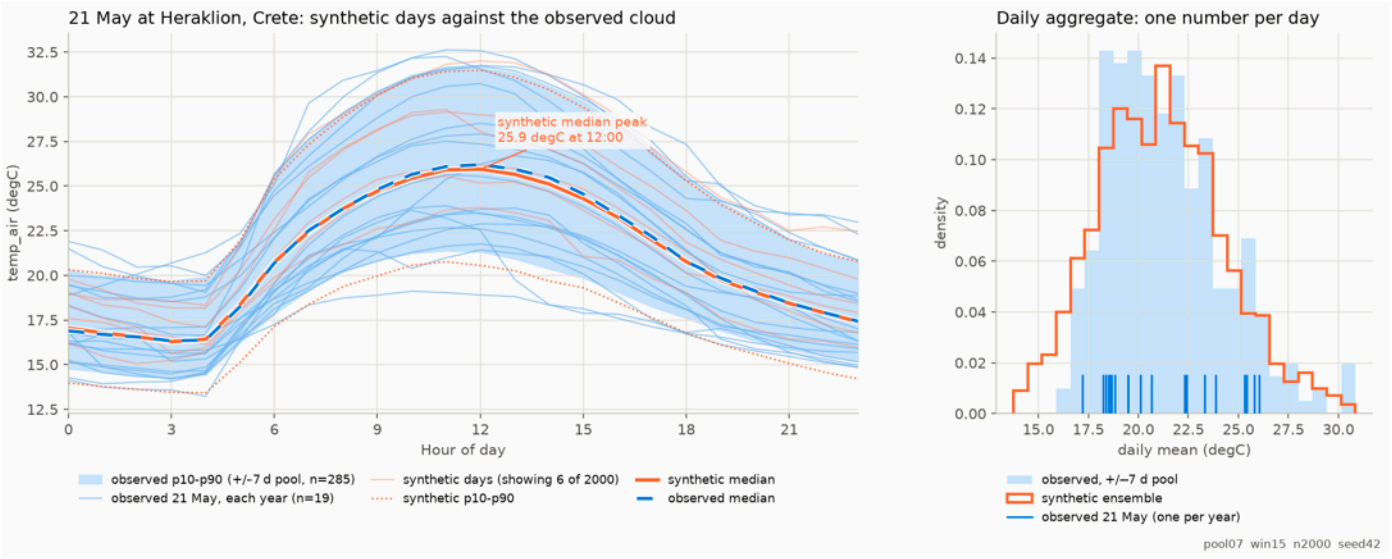


Figure 9. Typical day (21 May) at Heraklion, air temperature. (**Left**): hourly values. (**Right**): distribution of the daily mean.

4.3. Annual Marginal Fidelity

Marginal fidelity assesses whether the synthetic data take the right values, in the right proportions, regardless of their order in time. To prevent errors occurring in specific diurnal or seasonal regimes from being masked by aggregate annual averages (e.g., a single comparison over the whole year would mix noon with midnight and January with July, so large errors in one part of the day or year could be hidden by good agreement elsewhere), this fidelity is evaluated across each disaggregated month-by-hour cell, yielding 288 distinct bins (12×24). Within each cell, all historical observations across the 19-year record are compared against the corresponding synthetic sample across 1000 generated years using the first-order Wasserstein distance, W_1 [107]. The W_1 metric quantifies the average minimal transport displacement required to align synthetic and empirical distributions, expressed directly in physical units: W/m^2 for irradiance, °C for temperature, and m/s for wind speed.

Because historical weather records naturally exhibit year-to-year variability, cell distances are interpreted relative to empirical dispersion rather than an absolute zero threshold. An empirical benchmark threshold is constructed by randomly partitioning the 19-year Heraklion record into two subsets of 9 and 10 complete years and computing W_1 between them over 100 random splits. The 95th percentile of these historical-against-historical distances defines the benchmark threshold for each cell. For the identical 100 partitions, the first historical group is compared against a randomly drawn synthetic sample matched to the size of the second group; the median of these 100 resulting distances serves as the test statistic. A cell passes the marginal test when this median test statistic remains within the 95th percentile threshold, indicating that synthetic discrepancies do not exceed the natural dispersion observed between distinct historical year groupings. Nighttime solar irradiance cells—where observed and synthetic values are identically zero—are excluded from the evaluation pool (accounting for 139 cells under the ±7-day window configuration). In accordance with a decision rule fixed prior to analysis, a variable is considered statistically validated if at least 90% of the active cells pass this benchmark.

The results are presented as month-by-hour heatmaps, with the hour of day across and the month upwards. The left panel shows W_1 divided by the observed standard deviation of the cell, i.e., the mismatch in units of the cell's natural spread. This scaling makes cells comparable: a 4 W/m^2 difference is negligible at noon, when values range over hundreds of W/m^2, but large at sunrise, when they range over a few tens. The right panel shows the

difference between the synthetic and observed means, scaled the same way (red: synthetic higher, blue: lower). It separates a shift in the whole distribution from a difference in its width or shape. White crosses mark failed cells. Since each cell has its own threshold, the color shows how large a difference is, and the cross shows whether it exceeds what real years produce.

For irradiance (Figure 10), 145 of the 149 judged cells pass (97%), so the variable passes. Three failures lie at the edge of daylight (February 05:00, September 04:00, February 16:00), where irradiance depends strongly on the exact sunrise or sunset time, which shifts across the pooled days. The largest differences that still pass, such as November at 15:00, are also sunrise or sunset hours. The fourth failure, March at 04:00, is a special case: all observed values are zero, so the threshold is zero, and rare traces of synthetic light (mean 0.005 W/m^2) are enough to fail. The cell appears blank because its scaled distance cannot be computed (i.e., nighttime has zero irradiance with zero standard deviation). Through the main part of the day, the scaled distance is mostly below 0.05.

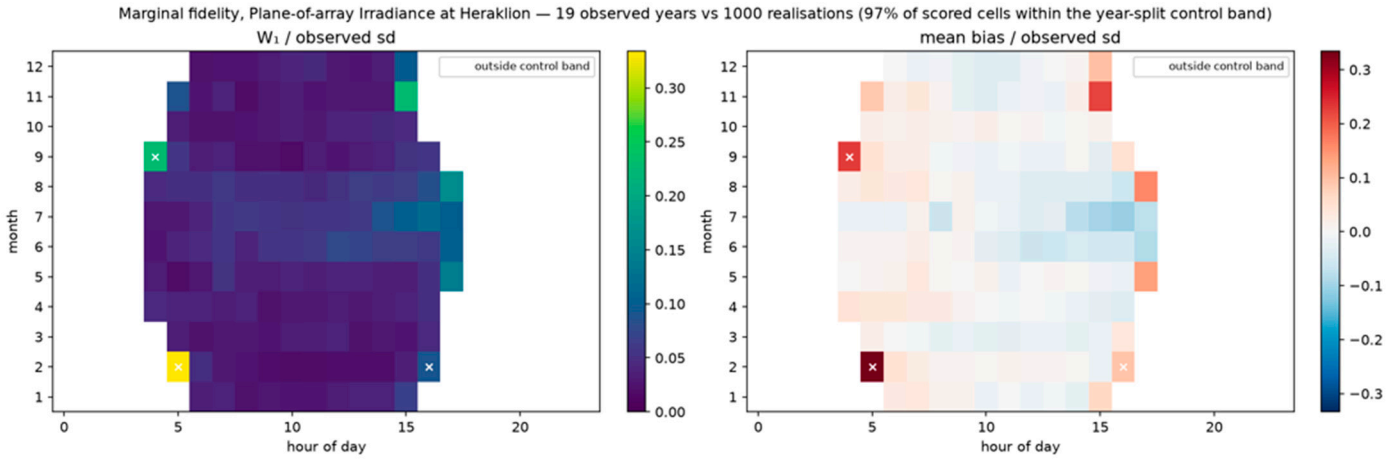


Figure 10. Marginal fidelity of plane-of-array irradiance, ±7-day ensemble (1000 realizations against 19 observed years). (**Left**): W_1 per cell, scaled by the observed standard deviation. (**Right**): difference in mean, scaled the same way. Crosses mark failed cells.

For air temperature and wind speed (Figures 11 and 12), all 288 cells pass. The typical mismatch is between a twentieth and an eighth of the natural spread (median scaled W_1 0.06 for wind speed and 0.13 for temperature).



Figure 11. Marginal fidelity of air temperature, ±7-day ensemble. Panels as in Figure 10.

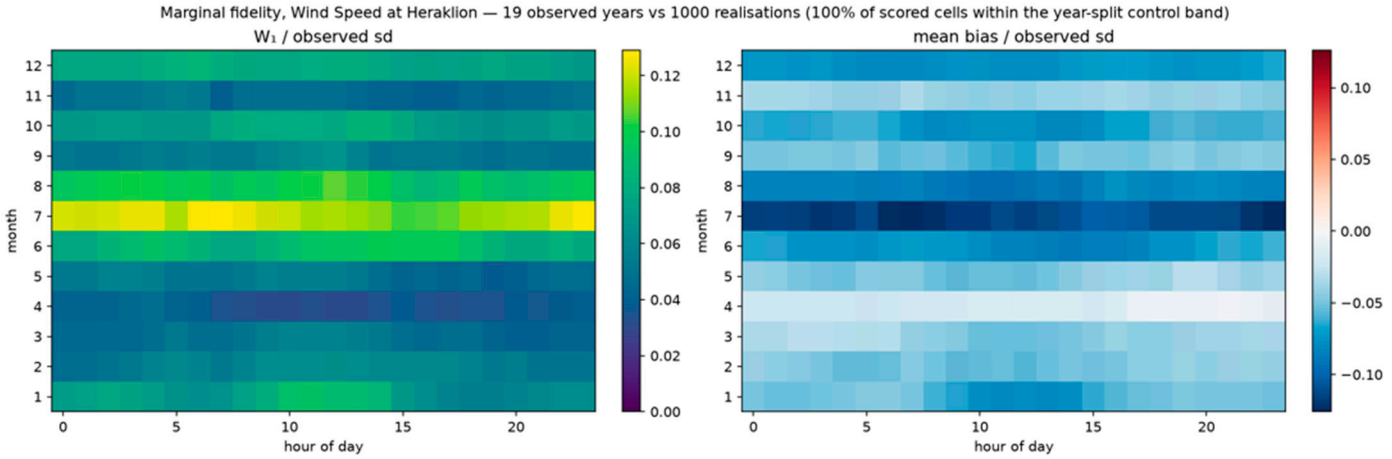


Figure 12. Marginal fidelity of wind speed, ±7-day ensemble. Panels as in Figure 10.

4.4. Annual Temporal Fidelity

While marginal tests evaluate distributional fidelity irrespective of temporal ordering—yielding identical results if hourly values were randomly shuffled—meteorological processes possess physical inertia and multi-scale temporal memory. For example, high afternoon temperatures naturally sustain warm evening regimes, and sustained wind regimes often span multiple consecutive days. Temporal fidelity evaluates whether the synthetic series preserve these memory characteristics, using the autocorrelation function (ACF) as the primary diagnostic tool. Fidelity is assessed at three time scales: hour-to-day persistence (Section 4.4.2), multi-day spells (Section 4.4.3), and year-to-year variability (Section 4.4.4); the dependence between variables is reported in Section 4.4.5. The metrics are defined in Section 4.4.1, and findings are discussed in Section 4.6.

4.4.1. Persistence Metrics

The autocorrelation at a discrete lag k measures the linear dependence between the time series X_t and its time-shifted counterpart X_{t+k}, ranging from 1 for identical linear agreement to 0 for absent linear relationship, with negative values reflecting anti-correlated fluctuations. For example, an autocorrelation of 0.8 at a 3 h lag indicates that an observed value carries substantial predictive power regarding the state three hours later. The ACF was computed across each continuous 8760 h series for all historical baseline years (M = 19) and synthetic realizations (N = 1000), evaluating lags up to k = 240 h (10 days). Individual years were evaluated independently rather than joined sequentially, preventing artificial step discontinuities at year boundaries from distorting long-lag estimates. To ensure smooth, monotonic decay profiles across extended lags, the standard sample autocovariance estimator was adopted [108].

Model persistence was examined using two complementary implementations of the autocorrelation function. The raw ACF evaluates unprocessed hourly series, capturing the combined effects of the diurnal cycle and annual seasonality. Because solar forcing imposes strict 24 h periodicity—where temperature at 14:00 correlates more strongly with the same hour the following day than with 02:00—the raw ACF naturally rebounds every 24 h, testing whether the generator reproduces cyclic diurnal and seasonal shapes. In contrast, the anomaly ACF is evaluated on residual series obtained by subtracting the baseline diurnal expectation for each calendar hour (calculated from the historical record using a smoothed window) from both observed and synthetic datasets. This detrending isolates departures from typical conditions, such as persistent heatwaves or wind spells, directly measuring a meteorological event duration [109]. Because the generator is formulated as a first-order

Markov chain conditioned solely on the immediately preceding hour, hourly transitions are matched by construction. The real test of temporal fidelity is therefore whether multi-day memory from synthetic data, which emerges only through sequential state transitions, matches empirical persistence.

Temporal memory was summarized using the integral decorrelation time τ, evaluated by summing the normalized autocorrelation $\rho(k)$ up to the first zero-crossing lag $k_{\max} = \min\{k \in \mathbb{N} \mid \rho(k) \leq 0\} - 1$, or truncated at 240 h if the autocorrelation remains strictly positive:

$$\tau = 1 + 2 \cdot \sum_{k=1}^{k_{max}} \rho(k)$$

The metric τ represents the effective persistence length of the underlying process, quantifying the number of correlated hourly increments required before successive observations behave as statistically independent samples [110,111]. Smaller values of τ indicate rapid atmospheric variability, whereas larger values reflect prolonged persistent weather regimes. For ambient temperature, the sample autocorrelation remains positive throughout the entire 240 h evaluation window, meaning that τ serves as a conservative lower bound; however, because the identical truncation horizon is applied to both empirical and synthetic series, the comparative assessment remains consistent and unbiased.

Throughout this section, τ is computed on the anomaly ACF, so that it measures how long departures from the diurnal–seasonal normal persist rather than the periodic cycle itself. It is also the more informative measure for comparison: the observed band in the figures spans single years, whereas the synthetic curve is the mean of 1000 realizations, so the averaged curve stays inside the band easily.

To translate persistence into quantities of practical relevance, three types of days are defined. A dull day has a daily irradiance total below half of the observed average for its month, a calm day has a daily mean wind speed below 70% of the monthly average, and a hot day has a daily mean temperature above the observed 90th percentile of its month. For each type, the probability that such a day is followed by another and the annual number of runs of three or more consecutive days are counted.

Year-to-year variability is quantified as the standard deviation, across years, of the annual irradiance total and of the annual mean temperature and wind speed. For the observed record, the linear trend over 2005–2023 is removed first, so that long-term drift is not counted as interannual variability.

4.4.2. Hour-to-Day Memory: ACF and Decorrelation Time

Figures 13–15 show the raw and anomaly ACF of each variable, and Table 2 summarizes the share of lags inside the observed band together with τ. A synthetic curve above the band means that departures last too long; below it, memory is lost too quickly. Ordered by the ratio of synthetic to observed τ, irradiance loses its memory too quickly ($\times 0.65$), air temperature is close to observed ($\times 1.21$), and wind speed persists too long ($\times 1.82$).

Irradiance. The raw ACF of the synthetic data lies on top of the observed one, so the day–night cycle is accurately reproduced. The anomaly ACF tells a different story. The historical curve band has bumps at 24 h and 48 h: a cloudy day tends to be followed by another cloudy day, because weather systems often last several days. The synthetic curve has no such bumps. Its memory is gone within about 10 h, and τ (5.3 h) equals the lowest value of any observed year.

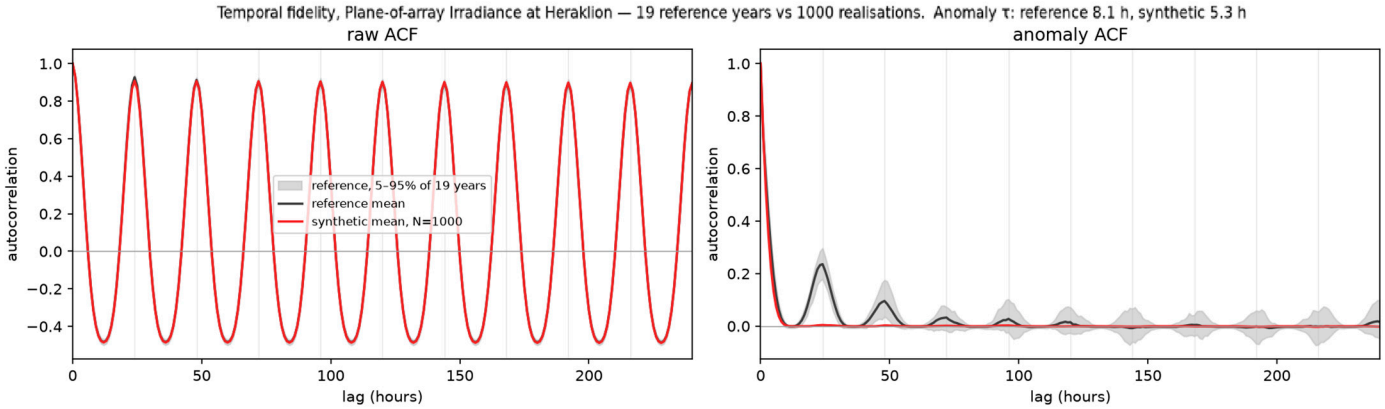


Figure 13. Temporal fidelity of plane-of-array irradiance, ±7-day ensemble. Left: raw ACF; right: anomaly ACF. Grey band: 5th–95th percentile of the 19 observed years; dark line: observed mean; red line: mean of 1000 synthetic years.

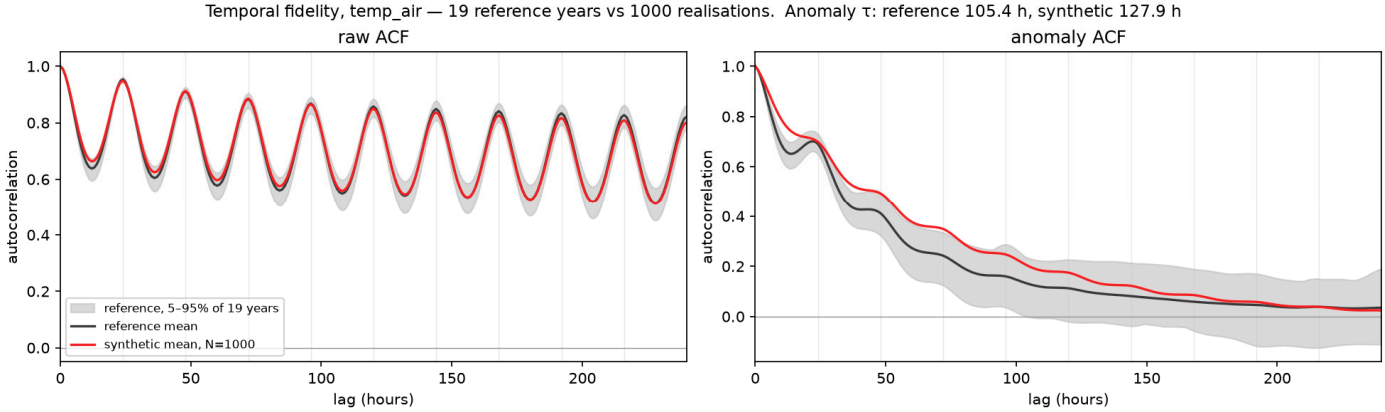


Figure 14. Temporal fidelity of air temperature, ±7-day ensemble. Panels as in Figure 13.

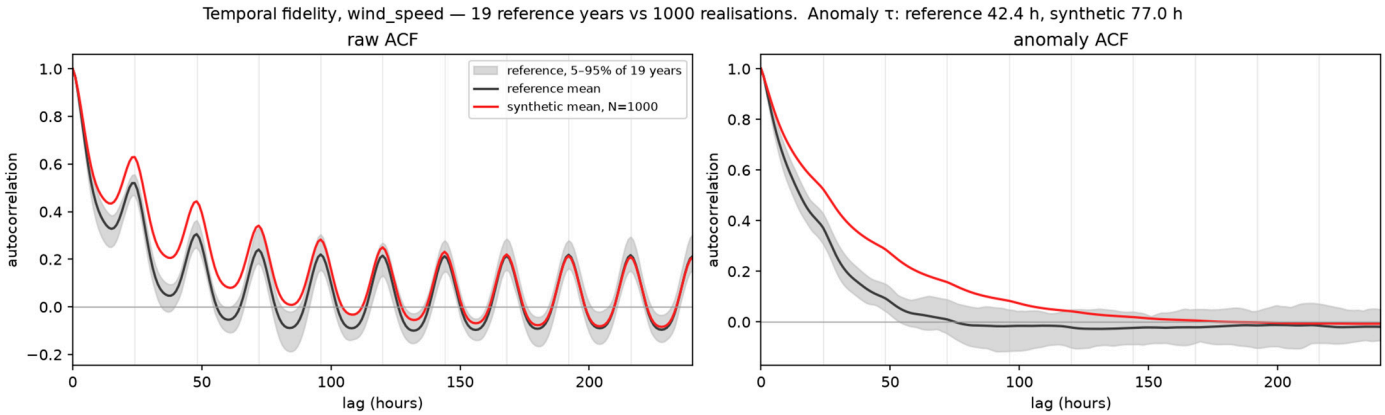


Figure 15. Temporal fidelity of wind speed, ±7-day ensemble. Panels as in Figure 13.

Air temperature. The raw ACF is inside the band at every lag. The anomaly ACF lies on the upper edge of the band up to about 100 h, so deviations from the seasonal normal persist slightly longer than observed: warm and cold spells last about 20% longer (τ 128 h against 105 h), which is still well within the range of real years (58–205 h).

Table 2. Temporal fidelity, ±7-day ensemble.

	Irradiance	Air Temperature	Wind Speed
Raw ACF, lags inside the band	98%	100%	60%
Anomaly ACF, lags inside the band	86%	75%	52%
τ observed, mean of historical data (h)	8.1	105	42
τ observed, range of historical data (h)	5.3–13.1	58–205	33–62
τ synthetic (h)	5.3	128	77
Ratio synthetic/observed	×0.65	×1.21	×1.82

Wind speed. The raw ACF is inside the band at only 60% of lags. All lags outside it, from 3 h to 107 h, lie above the band; after about 110 h the curve returns into it with daily ripples of the correct size. The day–night cycle is therefore reproduced, and what fails is that wind conditions persist too long. The anomaly ACF confirms this: it stays above the band up to about 120 h, and τ is 77 h against 42 h, longer than in any observed year (at most 62 h). Windy and calm spells therefore last almost twice as long as observed.

It can be seen that although the same generator is used, and that the raw ACF appears to perform as expected, the anomaly ACF results exhibit opposing behaviors, i.e., the irradiance shows reduced persistence, while wind and air temperature exhibit slightly longer persistence. This result is discussed in more detail in Section 4.6, alongside the rest of the findings.

4.4.3. Multi-Day Spells

Table 3 turns these findings into counts of the spells that matter in practice. In the observed data, a dull day is followed by another dull day 31% of the time; in the synthetic data, only 11% of the time, close to the 8% expected if days were fully independent. Runs of three or more dull days are therefore about six times rarer. Wind shows the opposite: calm days follow calm days more often than observed (56% against 44%) and runs of three or more calm days are about 45% more frequent. Hot spells are reproduced well, even though the temperature τ is somewhat long.

Table 3. Multi-day spells, ±7-day ensemble. Observed: mean of 19 years.

	Observed	Synthetic
Dull day followed by another dull day	31%	11%
Runs of ≥3 dull days per year	1.8	0.3
Calm day followed by another calm day	44%	56%
Runs of ≥3 calm days per year	9.3	13.5
Hot day followed by another hot day	56%	59%
Runs of ≥3 hot days per year	4.8	5.0

4.4.4. Interannual Variability

The same pattern reaches the scale of whole years (Table 4). The annual irradiance total varies about 40% less between synthetic years than between real years, whereas synthetic annual mean wind speed varies about 30% more. Annual mean temperature also varies less between synthetic years (0.34 °C against 0.47 °C), although its spells are, if anything, too long. These relative differences concern quantities that are themselves small compared with the annual values. Expressed as a coefficient of variation (CoV, standard deviation over the mean), the annual irradiance total varies by 1.1% between synthetic years against 1.7% between observed years, and the annual mean wind speed by 4.4% against 3.4%. For temperature, whose Celsius scale has no natural zero, the difference is best read in absolute terms: 0.13 °C.

Table 4. Year-to-year variation, ±7-day ensemble. Observed: standard deviation over 19 years after removing the linear trend over 2005–2023. In brackets: coefficient of variation relative to the observed annual mean (1879 kWh/m^2, 2.47 m/s).

Annual Year-to-Year Standard Deviation	Observed	Synthetic
Irradiance total	32 kWh/m^2 (1.7%)	20 kWh/m^2 (1.1%)
Mean temperature	0.47 °C	0.34 °C
Mean wind speed	0.085 m/s (3.4%)	0.109 m/s (4.4%)

For PV yield assessment, the narrower spread affects the P90 value. With a normal approximation (P90 $\approx \mu - 1.28\sigma$), the observed spread places the annual P90 about 41 kWh/m^2 (2.2%) below the mean, and the synthetic spread about 26 kWh/m^2 (1.4%) below it. The P50 is unaffected, and the P90 read from the raw ensemble differs from the observed one by about 0.8% of the annual yield at this site.

4.4.5. Dependence Between Variables

Because the variables are generated independently, the correlation between their daily departures from normal should vanish in the synthetic data. Table 5 measures this against the historic record. Daily anomalies are the daily irradiance total, mean temperature, and mean wind speed minus their observed calendar-day climatology (smoothed over ±7 days). In the observed record, all pairs are clearly correlated: sunny days are warmer and calmer, and warm days are calmer. The links are strongest in winter, when cloud, cool air, and strong wind arrive together with passing storms, and the sunshine–temperature link is weakest in summer, when skies are almost always clear. In the synthetic ensemble, every correlation is zero to within ±0.01, as expected from independent generation. A consequence is that unusual combinations are not excluded: a heavily overcast noon can coincide with an unusually high temperature. Such hours remain rare, because each condition is rare at noon on its own, but they are somewhat more frequent than in reality, where overcast noons tend to be cooler. Historical data for single years scatter widely: for the irradiance–temperature pair outside winter, and for the irradiance–wind pair in summer, the 5th–95th percentile range of single-year correlations includes zero. The observed dependence is clearest in the winter links of irradiance with temperature (0.27 to 0.51 in single years) and with wind (−0.56 to −0.27), and in the temperature–wind link from spring to autumn (for example, −0.70 to −0.35 in summer). These values set the target that joint generation (Section 6.3) will be tested against.

Table 5. Correlation of daily anomalies between variables, observed (19 years)/synthetic (±7-day ensemble, 1000 years). G: daily irradiance total; T: daily mean temperature; W: daily mean wind speed.

Season Months	Winter DJF	Spring MAM	Summer JJA	Autumn SON	Year All
G–T	0.39/0.00	0.22/0.00	0.13/0.00	0.24/0.01	0.24/0.00
G–W	−0.42/0.00	−0.21/0.00	0.01/0.00	−0.32/0.00	−0.25/0.00
T–W	−0.31/0.00	−0.35/−0.01	−0.53/0.00	−0.43/0.00	−0.39/0.00

4.5. Effect of Window Days on Fidelity

The training window is the main setting a user has to choose, and it involves a trade-off. Each calendar day's models are fitted to the days in its pool. A narrow window keeps the pool close to the target date, so the days in it share the same day length, sun height, and typical weather, but it gives the models fewer days to learn from. A wide window gives more days, but the pool then mixes days from different parts of the season. To see which

effect dominates, the whole evaluation of Sections 4.3 and 4.4 was repeated with a ±15-day window, and the results are compared in Table 6. In each row, the better of the two values indicates the better window for that measure: more passing cells, a smaller W_1, a τ ratio closer to 1, and more lags inside the band.

Table 6. Comparison of the ±7-day and ±15-day ensembles.

	±7 days	**±15 days**
Irradiance, cells passing the marginal test	97% (145/149), pass	88% (134/152), fail
Air temperature/wind speed, cells passing	100%/100%	100%/100%
Median scaled W_1: irradiance/temperature/wind	0.04/0.13/0.06	0.07/0.19/0.07
τ ratio: irradiance/temperature/wind	×0.65/×1.21/×1.82	×0.67/×1.32/×1.84
Anomaly ACF inside the band: irradiance/temperature/wind	86%/75%/52%	88%/64%/51%

Table 6 compares the two training windows. For irradiance (Figure 16), the ±15-day window raises the failed cells from 4 to 18, and irradiance falls below the 90% rule. Almost all failures lie at sunrise and sunset (04:00–05:00 and 15:00–17:00). This is expected: a ±15-day pool spans a month, during which sunrise and sunset move by more than half an hour around the equinoxes, so the values of a sunrise hour are mixed with days on which the sun was already higher or still below the horizon. The wider window also leaves traces of synthetic light in three more pre-dawn cells, which is why 152 cells are judged instead of 149.

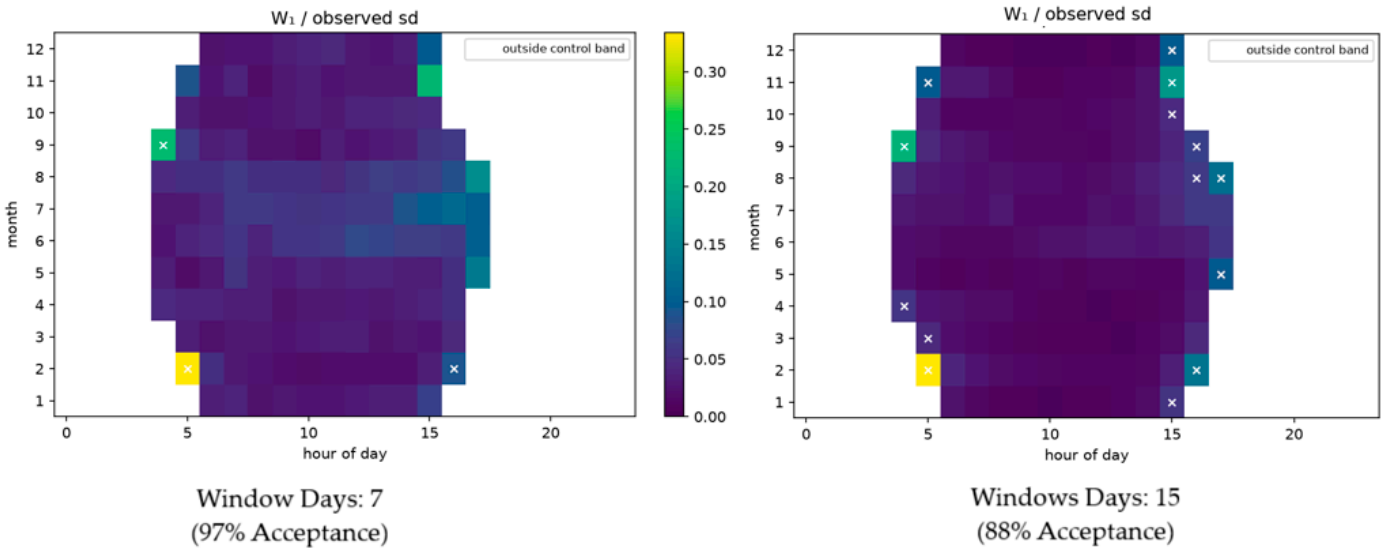


Figure 16. Marginal test for irradiance with a ±7-day (**left**) and a ±15-day (**right**) training window. Crosses mark failed cells.

The problem lies in pooling days by clock hour. If the days in a pool were instead lined up by the position of the sun, for example, by hours after sunrise, a sunrise hour would always be compared with other sunrise hours. A wider window could then be used without this side effect, which would help most where the record is short.

The wider window also affects air temperature: all cells still pass, but the typical mismatch grows by half (scaled W_1 0.13 to 0.19) and spells last longer (τ ratio 1.21 to 1.32). Wind speed barely changes.

Figure 17 shows the distribution of all hourly air temperatures of the year for the two windows. With all hours put together, observed and synthetic curves are almost identical in both cases; the differences above only appear when the year is split into month-by-hour cells. This is why the cell-by-cell test is needed rather than a single comparison for the whole year.

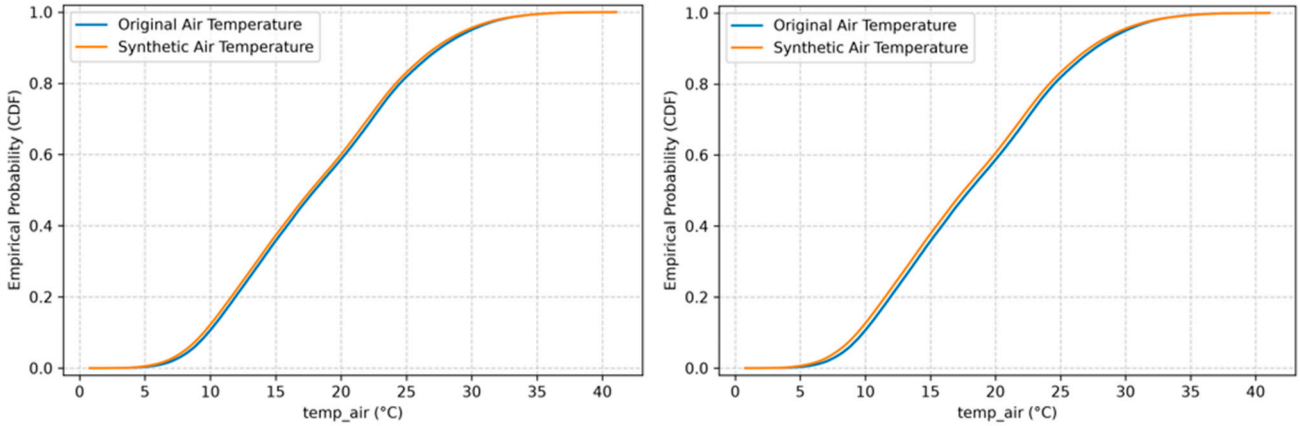


Figure 17. Cumulative distribution of hourly air temperature over the whole year, observed and synthetic, for the ±7-day (**left**) and ±15-day (**right**) training windows.

With 19 years of data, the ±7-day window is the better choice: it passes every marginal test and matches more closely on almost every measure. The benefit of a larger window is more data per model. With ±7 days and 19 years, each pool already holds 285 days, but with a shorter record the pools shrink, and a wider window may be needed. This trade-off was not exhaustively tested here and should be checked for each record length.

4.6. Possible Causes of the Discrepancies

Taken together, the findings of Sections 4.2–4.5 give a consistent picture. The generator reproduces the values the weather takes, hour by hour and month by month: almost all marginal cells pass, and the few failures sit at sunrise and sunset. It reproduces the daily and seasonal cycles. What the generator reproduces less well is how long departures from normal persist. Irradiance forgets a cloudy day too quickly, wind speed remembers too long, air temperature lies close to the observed range but on its long side, and the spread between whole years is too narrow for irradiance and temperature and too wide for wind. This section discusses what may cause these differences. The explanations below are hypotheses that fit the findings of this section; experimentally establishing the causes is outside the scope of this work.

4.6.1. The Night-Time Reset of Irradiance

The most likely explanation for the short memory of irradiance lies in how the chain passes information from one day to the next. Each hour only knows the value of the hour before it. For irradiance, every night is a "degenerate" stretch (Section 3.3.3) in which the value is exactly zero, whether the day before was clear or overcast. When the sun rises, the chain appears to start from the same zero state every morning, so each day begins with a blank slate and yesterday's cloudiness is lost. This would explain why the 24 h and 48 h bumps of the observed anomaly ACF are missing (Figure 13), and why a synthetic dull day is followed by another almost as rarely as if days were independent (11% against 8%, Table 3). Wind speed and air temperature never pass through such a stretch. Their value

at 23:00 differs from night to night, and the cross-midnight model carries it into the next morning, so a windy evening tends to lead to a windy morning.

4.6.2. First-Order Memory Applied over Many Days

The night-time reset explains why wind and temperature keep their memory across midnight, but not why wind keeps too much of it. One candidate is the first-order assumption itself. The chain reproduces the hour-to-hour dependence by construction (Section 4.4.1), and that dependence is very strong: the observed anomaly correlation between one hour and the next is about 0.99 for temperature and 0.97 for wind speed. A process that carries memory only from one hour to the next, and matches this correlation, keeps its departures for a long time. A simple check on the observed record suggests how long. A first-order autoregressive process with the observed hour-to-hour correlation has a decorrelation time of about 70 h for wind speed, against 42 h observed and 77 h synthetic, and about 140 h for temperature, against 105 h observed and 128 h synthetic. If this approximation holds, real wind and temperature lose their memory faster over several days than their hour-to-hour behavior alone would imply. A chain that sees only the previous hour cannot represent this, and most of the excess wind persistence could stem from this mechanism alone. The check is approximate, because the generator is neither linear nor Gaussian and its transitions change with the hour of day, so it indicates a plausible order of magnitude rather than a measured cause.

4.6.3. Pooling of the Training Window

A further contribution to the discrepancies in the temporal fidelity may rise from the way the training pool is built. Each pool mixes days from 19 different years and from up to two weeks either side of the target date. Days in the pool therefore differ in their background level: an early-spring day is cooler than a late-spring one, and one year may be windier than another. Within each day, the hours share that background level, so pairs of consecutive hours from the pool are correlated partly because they belong to the same background, not only because of weather persistence. The kernel cannot distinguish the two. During generation, the chain may then treat a difference in background as if it were a spell and carry it forward, which would add to the persistence of wind and temperature.

The window comparison of Section 4.5 is consistent with this for temperature. Widening the window from ± 7 to ± 15 days lengthens the temperature spells (τ ratio 1.21 to 1.32), as expected if a wider pool spans a larger part of the seasonal rise or fall in temperature. In the observed record, the hour-to-hour correlation of the pooled pairs also rises slightly with the wider window (0.9865 to 0.9876), and under the approximation of Section 4.6.2, a rise in this size is roughly enough to lengthen τ by the observed amount. For wind speed, whose seasonal change within a month is small, neither the pooled correlation nor the τ ratio changes with the window ($\times 1.82$ and $\times 1.84$). Pooling over a wider window therefore seems a plausible cause of the extra temperature persistence, but unlikely to be the main cause of the excess wind persistence. The mixing of different years is present at every window width, so the comparison of Section 4.5 cannot isolate its effect; in the observed record the pooled correlation is close to the correlation of the anomaly series (0.972 against 0.970 for wind), which suggests that this part is also small.

4.6.4. Year-to-Year Spread

The annual-scale results may follow from the same mechanisms. If synthetic dull days do not cluster, they average out over a year, which would narrow the spread of annual irradiance totals; the long synthetic wind spells would have the opposite effect. The reduced variation in annual mean temperature, despite spells that are slightly too long, suggests

that part of the difference between real years comes from slow influences, such as a whole season being warmer than usual, that a chain with an hour-to-hour memory cannot create.

4.6.5. Implications

For wind energy and storage studies, overly persistent wind regimes mean fewer but longer calm periods than observed. Storage sized against the synthetic ensemble would therefore face more demanding low-wind sequences than the historical record suggests, so sizing results are likely to be conservative; for natural ventilation studies, the same bias implies longer stretches of weak driving force. Conversely, the short irradiance memory means that multi-day solar droughts, and the storage they call for, are likely to be underestimated. For systems that combine several weather inputs, the variables are generated independently, so the correlations of Section 4.4.5 are not reproduced.

4.6.6. Generality of the Figures

All the figures of this section come from one site, one training record and one set of generator parameters: a relative bandwidth of h = 0.1 and a ±7-day window, with only the window width varied (Section 4.5). Their size depends on the local climate and on these settings. Examples are the 40% narrower annual spread of irradiance, the six times rarer runs of dull days and the twice as long wind spells. How persistent cloud and wind regimes are differs from place to place, and both the bandwidth and the window width control how much the training pool is smoothed. The mechanisms identified in Sections 4.6.1 and 4.6.2 are structural, so the direction of the irradiance and persistence effects is expected to hold at other sites, but their size is not. The Heraklion values should therefore be read as indicative orders of magnitude, not as general properties of the generator. A Global Sensitivity Analysis over sites and generator parameters, planned as follow-up work (Section 6.3), will establish which parameters matter most.

5. Applications in Energy System Studies

5.1. Overview

A single weather year, such as a TMY, gives a single answer. An ensemble of synthetic years gives a range of answers, and with it the chance of a bad outcome. This section outlines four uses of MeteoSynth ensembles in energy studies. For each one, it also notes what the results of Section 4 mean in practice: where the ensemble can be trusted as it is, and where care is needed. Table 7 gives an overview.

Table 7. Applications of MeteoSynth ensembles, and Section 4 implications.

Application	What the Ensemble Provides	What to Keep in Mind
Building energy (Section 5.2)	Distributions of heating and cooling demand, peak loads and overheating hours	Hot spells are well reproduced; temperature and irradiance are linked only through the daily and seasonal cycle
Control and storage (Section 5.3)	Many plausible years on which to test control strategies	Climatological, not a forecast; dull spells too rare, calm spells too long; variables generated independently
PV yield and batteries (Section 5.4)	P50/P90 yield, battery autonomy, clipping	Year-to-year spread of PV yield narrower (CoV 1.1% against 1.7%); runs of dull days too rare to be considered in autonomy sizing
Sensitivity analysis (Section 5.5)	Weather as an explicit uncertain input	-

5.2. Building Energy Performance via Monte Carlo Simulations

Building simulation tools such as EnergyPlus and TRNSYS are usually run with one weather file. Overheating risk, for example, the number of hours above 26 °C under CIBSE TM52, then becomes a single yes-or-no answer. Running the same building model once for each year of an ensemble of several hundred synthetic years gives full distributions instead: of annual cooling energy, of peak cooling load and of overheating hours. Equipment can then be sized for a chosen level of risk, for example, a peak cooling load that is not exceeded in 95% of years, rather than for one typical year.

Section 4 supports this use for air temperature. Every month-by-hour cell passes the marginal test, and hot spells occur as often and last as long as observed (Table 3). This matters because overheating builds up over several hot days in a row. Two limits should be kept in mind. First, temperature and irradiance are generated separately. They share the daily and seasonal cycle, but an unusually hot synthetic day is not more likely to be unusually sunny. Solar gains during a heatwave are therefore those of a typical day of that date, and peak cooling loads may be underestimated. At Heraklion this effect is small in summer, where the observed correlation between daily irradiance and temperature anomalies is only 0.13, but larger in winter (0.39, Table 5), where it affects heating demand: synthetic cloudy winter days are not colder than average, as real ones are. Second, annual mean temperature varies less between synthetic years than between real ones (0.34 °C against 0.47 °C, a difference of 0.13 °C), so the spread of annual energy use will be somewhat narrower than observed. Events within a year, such as hot spells, are represented more reliably than the differences between years.

5.3. Stochastic HVAC Control and Energy Storage Optimization

MeteoSynth produces climatological weather, not forecasts. Each synthetic year starts from the typical conditions of its start date, not from today's measured weather, so it cannot replace a forecast in day-ahead control. What it can do is test control strategies before they are used. A battery or heat-pump schedule tuned on one weather year may work well by chance. Running it over hundreds of synthetic years shows how often it performs above or below expectations, for example, how often the store runs empty or grid penalties are paid. The same ensembles can be used to tune rule-based controllers, or to train learning-based ones, on far more varied weather than the historical record contains. For a stochastic optimization that needs a set of plausible days for a given date, the day generator of Section 4.2 supplies such a set directly.

Two findings of Section 4 affect stress tests of storage. Runs of dull days are too rare (Table 3), so a store that relies on solar input will look safer than it is. Calm spells are too long, so a store that relies on wind will look less safe than it is. In both cases, the error comes from day-to-day memory, not from the hourly values, which pass the marginal test. For hybrid PV–wind systems, the variables are generated independently in the present release (Section 4.4.5).

5.4. Photovoltaic (PV) Reliability and Tail-Risk Sizing

Financial underwriting of utility-scale PV projects requires estimating P90 annual energy yield (the generation level exceeded with 90% probability). From a historical record of POA irradiance (G), MeteoSynth generates multi-year ensembles from which P50 and P90 yields can be read directly. For off-grid PV-battery systems, the same ensembles can in principle be used to stress-test battery autonomy during "solar droughts", i.e., runs of consecutive dull days.

Section 4 gives two points to consider in these uses. First, the year-to-year spread of the annual irradiance total is about 40% narrower than observed (standard deviation 20

against 32 kWh/m^2, Table 4). In relative terms, the annual total varies by 1.1% between synthetic years against 1.7% between observed ones. The P50 is unaffected, and the P90 read from the raw ensemble differs from the observed one by about 0.8% of the annual yield at Heraklion (Section 4.4.4). Where this margin matters, the ensemble spread of annual totals can be rescaled to the observed one; the inter-annual term planned in Section 6.3 will remove the need for this step.

Second, for off-grid design, runs of three or more dull days occur 0.3 times per year in the synthetic data against 1.8 times per year in the observed record (Table 3). The cause is the night-time reset of irradiance (Section 4.6.1): every night the irradiance chain passes through exactly zero, the night-time transitions are of the CONSTANT kind (Section 3.3.3), and the cloudiness of the previous day is lost. Battery autonomy sized on the raw ensemble alone will therefore tend to be optimistic, and it should be checked against the longest dull sequences of the observed record until the day-to-day memory of Section 6.3 is added. Uses that depend on hourly and daily values rather than on multi-day runs, such as inverter clipping or the distribution of daily yield, are much less affected, although very dark days are slightly less frequent than observed (Section 4.2).

5.5. Global Sensitivity Analysis (GSA)

In Global Sensitivity Analysis (e.g., Variance-Based Sobol Analysis), researchers evaluate the relative contribution of design inputs (e.g., wall insulation, window solar heat gain coefficient, infiltration rate) to overall building performance variance. When weather is treated as a fixed deterministic constant (TMY), the sensitivity indices of envelope parameters are artificially inflated because meteorological variance is suppressed. By introducing the MeteoSynth ensemble, the weather becomes one more uncertain input: the realization index is sampled like any other parameter. Each realization is drawn from its own random stream, set only by the random seed, the realization number, and the variable, so a given index always gives the same weather year, whatever the size of the ensemble. The end effect is that a MeteoSynth ensemble acts as an explicit stochastic input variable in Sobol sampling schemes, giving researchers the tools to quantify the true main-effect (S_i) and total-effect (ST_i) indices of physical building parameters against natural microclimatic noise.

6. Conclusions and Future Roadmap

6.1. Summary of Contributions

MeteoSynth is an open-source Python package that generates continuous hourly synthetic weather series, and large ensembles of them, from historical records without assuming a parametric distribution. Hourly transitions follow a time-inhomogeneous first-order Markov chain whose densities are estimated by two-dimensional Kernel Density Estimation. Physical bounds are enforced by construction, zero-variance night-time irradiance is handled without numerical failure, a circular 366-day rolling window removes the New Year discontinuity and handles 29 February, and a cross-midnight model fitted on real consecutive dates links successive days. Fitting each day's models once for the whole ensemble makes 1000 synthetic years approximately two orders of magnitude faster to produce than generating them one by one. The package is variable-agnostic; in its present release it is a parallel univariate generator.

In the Heraklion case study (19 years of PVGIS data, ±7-day window), fidelity was judged against the spread between real years, cell by month-by-hour cell. All temperature and wind cells and 97% of active irradiance cells pass the marginal test; the remaining failures lie at sunrise and sunset, and the daily and seasonal cycles and hot spells are reproduced. The remaining differences lie in how long departures from normal last. Irradiance loses its day-to-day memory at night, so runs of dull days are about six times too rare, and the

year-to-year spread of annual irradiance is about 40% narrower than observed, although in relative terms this is a coefficient of variation of 1.1% against 1.7%; wind spells last about twice as long as observed. Because the variables are generated independently, their observed correlations (up to 0.5 in absolute value) are not reproduced. A ±7-day window outperformed a ±15-day one for this record length. These figures refer to one site, chosen for convenience, and one parameter set. They indicate orders of magnitude rather than general properties of the generator (Section 4.6.6).

The package is therefore well suited to studies driven by hourly, diurnal and seasonal statistics, such as building thermal performance, overheating, PV yield (P50) and sensitivity analysis. For applications driven by multi-day persistence, such as off-grid battery autonomy, or by the joint behavior of several variables, users should take the above differences into account.

6.2. Code Availability and Reproducibility

MeteoSynth (version 0.3.3) is available under the MIT license on GitHub (Supplementary Materials https://github.com/npapnet/meteosynth accessed on 21 November 2026) and on PyPI (pip install meteosynth). It requires Python ≥ 3.10 with NumPy ≥ 1.22, pandas ≥ 1.4, SciPy ≥ 1.8, pvlib ≥ 0.9, pyarrow, matplotlib, seaborn, and tqdm; exact lower bounds are declared in the package metadata, and a lock file is provided for the development environment. Documentation, with installation instructions, minimal examples, and worked use cases, is Supplementary Materials available at https://meteosynth-docs.npapnet-cloudflare.workers.dev/ (21 November 2026). The ensembles of Section 4 can be regenerated with a few lines:

```
from meteosynth import write_ensemble
from meteosynth.adapters import fetch_pvgis

ds = fetch_pvgis(35.3387, 25.1442, site="heraklion")  #
PVGIS 2005-2023, cached
write_ensemble(ds, attrs=["poa_global", "temp_air",
"wind_speed"],
path="ensemble.parquet", start=(1, 1), n_days=365,
n_realisations=1000, seed=2026, n_days_window=7)
```

The repository contains the scripts that produce every result and figure of this paper in its examples directory. Ensembles are stored in Parquet format together with the parameters that produced them and the provenance of the training record. Because each realization depends only on the random seed, the realization index, and the variable, any realization can be regenerated independently.

6.3. Future Development Roadmap

The results point to a clear order of work.

Day-to-day memory of irradiance. A latent signal must carry cloudiness across the night, for example, by conditioning the morning transitions on the previous day's insolation, or by a regime-switching mechanism between clear and overcast multi-day states. This would restore the frequency of dull spells, which matters most for off-grid storage sizing.

Inter-annual variability. An hour-to-hour chain trained on a pooled multi-year record cannot create a whole season that is warmer or cloudier than usual. Superimposing a low-frequency seasonal or annual perturbation, drawn once per realization, would widen

the ensemble spread of annual totals towards the observed one. Together with the previous item, this is the priority for PV yield (P90) applications.

Joint multivariate generation. This is the key next step of the package: Section 4.4.5 shows that the observed correlations between the daily departures of the variables are not reproduced by independent generation. Joint transition models, through multidimensional Kernel Density Estimation or pair-copula constructions, will be introduced once the single-variable memory above is in place, so that the coupling is not fitted to a chain whose own persistence is still biased. The validation in Section 4.4.5 will serve as its acceptance test.

Sensitivity to site and parameters. A follow-up study will apply Global Sensitivity Analysis to the generator itself, over sites of different climate and over its parameters (kernel bandwidth, window width and record length). It will show how far the Heraklion figures generalize, and which parameters should be tuned first.

Further refinements. The excess persistence of wind speed, and to a lesser degree of temperature, needs calibration (Section 4.6). Aligning pooled days by solar position rather than clock hour would allow wider windows without the sunrise and sunset errors of Section 4.5. Connectors to data sources other than PVGIS, such as ERA5 directly or ground-station records, will support validation across several sources and sites. Finally, climate-change scenarios could be introduced by shifting the training pools towards projected future distributions.

Supplementary Materials: The following supporting information can be downloaded at: the PyPI package is at https://pypi.org/project/meteosynth/ accessed on 21 November 2026; the public GitHub repository (MIT license), including the scripts that produce every result of this paper, is at https://github.com/npapnet/meteosynth. The documentation, with installation instructions, minimal reproducible examples, and worked use cases, is at https://meteosynth-docs.npapnet-cloudflare.workers.dev/.

Funding: This research received no external funding.

Data Availability Statement: The data supporting this work can be reproduced by executing the scripts of the public repository, which download the PVGIS record for the site and regenerate the ensembles from fixed random seeds.

Acknowledgments: During the preparation of this manuscript/study, the author used Anthropic Claude 5.5 for language editing, drafting, and revising text, and assisting with benchmarking analysis scripts, and Google Gemini for creating Figures 1–4 and 6. The author has reviewed and edited the output and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABC-GAN	Attention-Based Cycle-Consistent Generative Adversarial Network
ACF	autocorrelation function
API	Application programming interface
ARMA	Autoregressive Moving Average
CTG-CSA	Conditional Dual-Attention Generative Adversarial Framework
CDF	Cumulative Probability Distribution
CF	Climate and Forecast (metadata conventions)
CNN	Convolutional Neural Network
CSDI	Conditional Score-based Diffusion Imputation
CVAE	Conditional Variational Autoencoders
DCGAN	Deep Convolutional GANs
DDPM	Denoising Diffusion Probabilistic Models

DFT	Discrete Fourier Transform
DGM	Deep Generative Models
DNI	Direct Normal Irradiance
ECDF	Empirical Cumulative Distribution Function
EENS	Expected Energy Not Served
ELBO	Evidence Lower Bound
FARIMA	Fractional Autoregressive Integrated Moving Average
GAN	Generative Adversarial Networks
GARCH	Generalized Autoregressive Conditional Heteroskedasticity
GHI	Global Horizontal Irradiance
GPU	Graphics Processing Unit
GSA	Global Sensitivity Analysis
IWEC	International Weather for Energy Calculations
KDE	Kernel Density Estimation
KDE 2D	Two-Dimensional Kernel Density Estimation
LOLE	Loss of Load Expectation
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MBB	Moving Block Bootstrap
MTM	Markov Transition Matrices
NWP	Numerical Weather Prediction
PCC	Pair-Copula Constructions
POA	Plane of Array
PVGIS	Photovoltaic Geographical Information System
RNN	Recurrent neural networks
RMSE	Root Mean Square Error
SOM	Self-Organizing Map
SRAI	Statistical Relational AI
STL	Seasonal and Trend decomposition using Loess
STSG	Synthetic Time Series Generation
TAG	Time-dependent Autoregressive Gaussian
TCN	Temporal Convolutional Network
TMY	Typical Meteorological Years
VAE	Variational Autoencoders
VAR	Vector Autoregressive
VRE	Variable Renewable Energy
WGAN-GP	Wasserstein GAN with Gradient Penalty
WGEN	Richardson-type Weather Generator

References

1. Šúri, M.; Huld, T.A.; Dunlop, E.D.; Ossenbrink, H.A. Potential of Solar Electricity Generation in the European Union Member States and Candidate Countries. *Sol. Energy* **2007**, *81*, 1295–1305. [CrossRef]
2. Anderson, K.S.; Hansen, C.W.; Holmgren, W.F.; Jensen, A.R.; Mikofski, M.A.; Driesse, A. Pvlib Python: 2023 Project Update. *J. Open Source Softw.* **2023**, *8*, 5994. [CrossRef]
3. Jensen, A.R.; Anderson, K.S.; Holmgren, W.F.; Mikofski, M.A.; Hansen, C.W.; Boeman, L.J.; Loonen, R. Pvlib Iotools—Open-Source Python Functions for Seamless Access to Solar Irradiance Data. *Sol. Energy* **2023**, *266*, 112092. [CrossRef]
4. Holmgren, W.F.; Hansen, C.W.; Mikofski, M.A. Pvlib Python: A Python Package for Modeling Solar Energy Systems. *J. Open Source Softw.* **2018**, *3*, 884. [CrossRef]
5. Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. The ERA5 Global Reanalysis. *Q. J. R. Meteorol. Soc.* **2020**, *146*, 1999–2049. [CrossRef]
6. Jiao, Z.; Yuan, J.; Farnham, C.; Emura, K. Multivariate Stochastic Generation of Meteorological Data for Building Simulation through Interdependent Meteorological Processes. *Sci. Rep.* **2024**, *14*, 24927. [CrossRef] [PubMed]
7. Hong, T.; Pinson, P.; Fan, S.; Zareipour, H.; Troccoli, A.; Hyndman, R.J. Probabilistic Energy Forecasting: Global Energy Forecasting Competition 2014 and Beyond. *Int. J. Forecast.* **2016**, *32*, 896–913. [CrossRef]

8. Bessa, R.; Möhrlen, C.; Fundel, V.; Siefert, M.; Browell, J.; Haglund El Gaidi, S.; Hodge, B.-M.; Cali, U.; Kariniotakis, G. Towards Improved Understanding of the Applicability of Uncertainty Forecasts in the Electric Power Industry. *Energies* **2017**, *10*, 1402. [CrossRef]
9. Wang, J.; Botterud, A.; Bessa, R.; Keko, H.; Carvalho, L.; Issicaba, D.; Sumaili, J.; Miranda, V. Wind Power Forecasting Uncertainty and Unit Commitment. *Appl. Energy* **2011**, *88*, 4014–4023. [CrossRef]
10. Zheng, Q.P.; Wang, J.; Liu, A.L. Stochastic Optimization for Unit Commitment—A Review. *IEEE Trans. Power Syst.* **2015**, *30*, 1913–1924. [CrossRef]
11. Ortega-Vazquez, M.A.; Kirschen, D.S. Estimating the Spinning Reserve Requirements in Systems With Significant Wind Power Generation Penetration. *IEEE Trans. Power Syst.* **2009**, *24*, 114–124. [CrossRef]
12. Matos, M.A.; Bessa, R.J. Setting the Operating Reserve Using Probabilistic Wind Power Forecasts. *IEEE Trans. Power Syst.* **2011**, *26*, 594–603. [CrossRef]
13. Sweeney, C.; Bessa, R.J.; Browell, J.; Pinson, P. The Future of Forecasting for Renewable Energy. *WIREs Energy Environ.* **2020**, *9*, e365. [CrossRef]
14. Pfenninger, S.; Hawkes, A.; Keirstead, J. Energy Systems Modeling for Twenty-First Century Energy Challenges. *Renew. Sustain. Energy Rev.* **2014**, *33*, 74–86. [CrossRef]
15. Bloomfield, H.C.; Gonzalez, P.L.M.; Lundquist, J.K.; Stoop, L.P.; Browell, J.; Dargaville, R.; Felice, M.D.; Gruber, K.; Hilbers, A.; Kies, A.; et al. The Importance of Weather and Climate to Energy Systems: A Workshop on Next Generation Challenges in Energy–Climate Modeling. *Bull. Am. Meteorol. Soc.* **2021**, *102*, E159–E167. [CrossRef]
16. Craig, M.T.; Wohland, J.; Stoop, L.P.; Kies, A.; Pickering, B.; Bloomfield, H.C.; Browell, J.; Felice, M.D.; Dent, C.J.; Deroubaix, A.; et al. Overcoming the Disconnect between Energy System and Climate Modeling. *Joule* **2022**, *6*, 1405–1417. [CrossRef]
17. Wilczak, J.M.; Akish, E.; Capotondi, A.; Compo, G.P.; Hoell, A. A Multi-Decadal Analysis of U.S. and Canadian Wind and Solar Energy Droughts. *J. Renew. Sustain. Energy* **2024**, *16*, 056502. [CrossRef]
18. Palma-Behnke, R.; Valencia, F.; Vega-Herrera, J.; Núñez-Mata, O. Synthetic Time Series Generation Model for Analysis of Power System Operation and Expansion with High Renewable Energy Penetration. *J. Mod. Power Syst. Clean Energy* **2021**, *9*, 849–858. [CrossRef]
19. Sakki, G.K.; Tsoukalas, I.; Kossieris, P.; Makropoulos, C.; Efstratiadis, A. Stochastic Simulation-Optimization Framework for the Design and Assessment of Renewable Energy Systems under Uncertainty. *Renew. Sustain. Energy Rev.* **2022**, *168*, 112886. [CrossRef]
20. Claassen, J.N.; Koks, E.E.; de Ruiter, M.C.; Ward, P.J.; Jäger, W.S. A Synthetic European Weather Dataset Based on Spatiotemporal Vine Copulas. *Sci. Data* **2025**, *12*, 1734. [CrossRef] [PubMed]
21. Turowski, M.; Heidrich, B.; Weingärtner, L.; Springer, L.; Phipps, K.; Schäfer, B.; Mikut, R.; Hagenmeyer, V. Generating Synthetic Energy Time Series: A Review. *Renew. Sustain. Energy Rev.* **2024**, *206*, 114842. [CrossRef]
22. Hall, I.J.; Prairie, R.R.; Anderson, H.E.; Boes, E.C. *Generation of a Typical Meteorological Year*; Sandia Labs.: Albuquerque, NM, USA, 1978.
23. Wilcox, S.; Marion, W. *Users Manual for TMY3 Data Sets*; Technical Report; National Renewable Energy Laboratory (NREL): Golden, CO, USA, 2008.
24. Petrakis, M.; Kambezidis, H.D.; Lykoudis, S.; Adamopoulos, A.D.; Kassomenos, P.; Michaelides, I.M.; Kalogirou, S.A.; Roditis, G.; Chrysis, I.; Hadjigianni, A. Generation of a "Typical Meteorological Year" for Nicosia, Cyprus. *Renew. Energy* **1998**, *13*, 381–388. [CrossRef]
25. Rahman, I.A.; Dewsbury, J. Selection of Typical Weather Data (Test Reference Years) for Subang, Malaysia. *Build. Environ.* **2007**, *42*, 3636–3641. [CrossRef]
26. Staffell, I.; Pfenninger, S. Using Bias-Corrected Reanalysis to Simulate Current and Future Wind Power Output. *Energy* **2016**, *114*, 1224–1239. [CrossRef]
27. Pfenninger, S.; Staffell, I. Long-Term Patterns of European PV Output Using 30 Years of Validated Hourly Reanalysis and Satellite Data. *Energy* **2016**, *114*, 1251–1265. [CrossRef]
28. van der Wiel, K.; Stoop, L.P.; van Zuijlen, B.R.H.; Blackport, R.; van den Broek, M.A.; Selten, F.M. Meteorological Conditions Leading to Extreme Low Variable Renewable Energy Production and Extreme High Energy Shortfall. *Renew. Sustain. Energy Rev.* **2019**, *111*, 261–275. [CrossRef]
29. Ruhnau, O.; Qvist, S. Storage Requirements in a 100% Renewable Electricity System: Extreme Events and Inter-Annual Variability. *Environ. Res. Lett.* **2022**, *17*, 044018. [CrossRef]
30. Calabrese, G. Generating Reserve Capacity Determined by the Probability Method. *Trans. Am. Inst. Electr. Eng.* **1947**, *66*, 1439–1450. [CrossRef]
31. Billinton, R.; Allan, R.N. *Reliability Evaluation of Power Systems*; Springer: New York, NY, USA, 1996; ISBN 978-0-306-45259-8.
32. Saket, R.K.; Padmanaban, S. *Reliability Analysis of Modern Power Systems*; IEEE eBooks; IEEE Xplore: Piscataway, NJ, USA, 2024; ISBN 978-1-394-22674-0.

33. Dent, C.J.; Keane, A.; Bialek, J.W. Simplified Methods for Renewable Generation Capacity Credit Calculation: A Critical Review. In Proceedings of the IEEE PES General Meeting, Minneapolis, MN, USA, 25–29 July 2010; pp. 1–8.
34. Jorgenson, J.; Awara, S.; Stephen, G.; Mai, T. *Comparing Capacity Credit Calculations for Wind: A Case Study in Texas*; National Renewable Energy Laboratory (NREL): Golden, CO, USA, 2021.
35. Gelaro, R.; McCarty, W.; Suárez, M.J.; Todling, R.; Molod, A.; Takacs, L.; Randles, C.A.; Darmenov, A.; Bosilovich, M.G.; Reichle, R.; et al. The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2). *J. Clim.* **2017**, *30*, 5419–5454. [CrossRef] [PubMed]
36. Liu, B.Y.H.; Jordan, R.C. The Interrelationship and Characteristic Distribution of Direct, Diffuse and Total Solar Radiation. *Sol. Energy* **1960**, *4*, 1–19. [CrossRef]
37. Richardson, C.W. Stochastic Simulation of Daily Precipitation, Temperature, and Solar Radiation. *Water Resour. Res.* **1981**, *17*, 182–190. [CrossRef]
38. Brown, B.G.; Katz, R.W.; Murphy, A.H. Time Series Models to Simulate and Forecast Wind Speed and Wind Power. *J. Appl. Meteorol. Climatol.* **1984**, *23*, 1184–1195. [CrossRef]
39. Wilks, D.S.; Wilby, R.L. The Weather Generation Game: A Review of Stochastic Weather Models. *Prog. Phys. Geogr. Earth Environ.* **1999**, *23*, 329–357. [CrossRef]
40. Ailliot, P.; Boutigny, M.; Koutroulis, E.; Malisovas, A.; Monbet, V. Stochastic Weather Generator for the Design and Reliability Evaluation of Desalination Systems with Renewable Energy Sources. *Renew. Energy* **2020**, *158*, 541–553. [CrossRef]
41. Graham, V.A.; Hollands, K.G.T.; Unny, T.E. A Time Series Model for *K*t with Application to Global Synthetic Weather Generation. *Sol. Energy* **1988**, *40*, 83–92. [CrossRef]
42. Aguiar, R.; Collares-Pereira, M. TAG: A Time-Dependent, Autoregressive, Gaussian Model for Generating Synthetic Hourly Radiation. *Sol. Energy* **1992**, *49*, 167–174. [CrossRef]
43. Zheng, K.; Sun, Z.; Song, Y.; Zhang, C.; Zhang, C.; Chang, F.; Yang, D.; Fu, X. Stochastic Scenario Generation Methods for Uncertainty in Wind and Photovoltaic Power Outputs: A Comprehensive Review. *Energies* **2025**, *18*, 503. [CrossRef]
44. Mora-López, L.; Sidrach-de-Cardona, M. Multiplicative ARMA Models to Generate Hourly Series of Global Irradiation. *Sol. Energy* **1998**, *63*, 283–291. [CrossRef]
45. Hasslet, J.; Raftery, A.E. Space-Time Modelling with Long-Memory Dependence: Assessing Ireland's Wind Power Resource. *J. R. Stat. Soc. Ser. C Appl. Stat.* **1989**, *38*, 1–50. [CrossRef]
46. Pesch, T.; Schröders, S.; Allelein, H.J.; Hake, J.F. A New Markov-Chain-Related Statistical Approach for Modelling Synthetic Wind Power Time Series. *New J. Phys.* **2015**, *17*, 055001. [CrossRef]
47. Lei, M.; Shiyan, L.; Chuanwen, J.; Hongling, L.; Yan, Z. A Review on the Forecasting of Wind Speed and Generated Power. *Renew. Sustain. Energy Rev.* **2009**, *13*, 915–920. [CrossRef]
48. Soman, S.S.; Zareipour, H.; Malik, O.; Mandal, P. A Review of Wind Power and Wind Speed Forecasting Methods with Different Time Horizons. In *Proceedings of the North American Power Symposium 2010*; IEEE: New York, NY, USA, 2010; pp. 1–8.
49. Wang, X.; Guo, P.; Huang, X. A Review of Wind Power Forecasting Models. *Energy Procedia* **2011**, *12*, 770–778. [CrossRef]
50. Foley, A.M.; Leahy, P.G.; Marvuglia, A.; McKeogh, E.J. Current Methods and Advances in Forecasting of Wind Power Generation. *Renew. Energy* **2012**, *37*, 1–8. [CrossRef]
51. Campbell, S.D.; Diebold, F.X. Weather Forecasting for Weather Derivatives. *J. Am. Stat. Assoc.* **2005**, *100*, 6–16. [CrossRef]
52. Richardson, C.W.; Wright, D.A. *WGEN: A Model for Generating Daily Weather Variables*; U.S. Dept. of Agriculture, Agricultural Research Service; Springfield, Va.; National Technical Information Service (NTIS): Washington, DC, USA, 1984.
53. Semenov, M.; Barrow, E.M. Use of a Stochastic Weather Generator in the Development of Climate Change Scenarios. *Clim. Change* **1997**, *35*, 397–414. [CrossRef]
54. Ohba, M.; Kanno, Y.; Nohara, D. Climatology of Dark Doldrums in Japan. *Renew. Sustain. Energy Rev.* **2022**, *155*, 111927. [CrossRef]
55. Cover, T.; Hart, P. Nearest Neighbor Pattern Classification. *IEEE Trans. Inf. Theory* **1967**, *13*, 21–27. [CrossRef]
56. Yakowitz, S. Nearest-Neighbour Methods for Time Series Analysis. *J. Time Ser. Anal.* **1987**, *8*, 235–247. [CrossRef]
57. Lall, U.; Sharma, A. A Nearest Neighbor Bootstrap For Resampling Hydrologic Time Series. *Water Resour. Res.* **1996**, *32*, 679–693. [CrossRef]
58. Rajagopalan, B.; Lall, U. A K-Nearest-Neighbor Simulator for Daily Precipitation and Other Weather Variables. *Water Resour. Res.* **1999**, *35*, 3089–3101. [CrossRef]
59. Amonkar, Y.; Farnham, D.J.; Lall, U. A K-Nearest Neighbor Space-Time Simulator with Applications to Large-Scale Wind and Solar Power Modeling. *Patterns* **2022**, *3*, 100454. [CrossRef] [PubMed]
60. Pavlides, A.; Agou, V.D.; Hristopulos, D.T. Non-Parametric Kernel-Based Estimation and Simulation of Precipitation Amount. *J. Hydrol.* **2022**, *612*, 127988. [CrossRef]
61. Gai, S.; Zhang, X.; Xie, J.; Liu, J. Kernel Density Estimation Method to Determine Coincident Design Day for Air-Conditioning System Design in Marine Climate of China. *Build. Environ.* **2025**, *269*, 112487. [CrossRef]

62. Smith, K.; Strong, C.; Rassoul-Agha, F. A New Method for Generating Stochastic Simulations of Daily Air Temperature for Use in Weather Generators. *J. Appl. Meteorol. Climatol.* **2017**, *56*, 953–963. [CrossRef]
63. Kunsch, H.R. The Jackknife and the Bootstrap for General Stationary Observations. *Ann. Stat.* **1989**, *17*, 1217–1241. [CrossRef]
64. Alonso, A.M.; Peña, D.; Romo, J. Forecasting Time Series with Sieve Bootstrap. *J. Stat. Plan. Inference* **2002**, *100*, 1–11. [CrossRef]
65. Wang, J.; Lin, Z.J.; Chen, W. Robust Predictions with Ambiguous Time Delays: A Bootstrap Strategy. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*; Association for Computing Machinery: New York, NY, USA, 2024; pp. 3104–3115.
66. Politis, D.N.; Romano, J.P. The Stationary Bootstrap. *J. Am. Stat. Assoc.* **1994**, *89*, 1303–1313. [CrossRef]
67. Aloni, O.; Perelman, G.; Fishbain, B. Synthetic Random Environmental Time Series Generation with Similarity Control, Preserving Original Signal's Statistical Characteristics. *Environ. Model. Softw.* **2025**, *185*, 106283. [CrossRef]
68. Sklar, M. Fonctions de Répartition à N Dimensions et Leurs Marges. *Ann. ISUP* **1959**, *VIII*, 229–231.
69. Papaefthymiou, G.; Kurowicka, D. Using Copulas for Modeling Stochastic Dependence in Power System Uncertainty Analysis. *IEEE Trans. Power Syst.* **2009**, *24*, 40–49. [CrossRef]
70. Bedford, T.; Cooke, R.M. Vines—A New Graphical Model for Dependent Random Variables. *Ann. Stat.* **2002**, *30*, 1031–1068. [CrossRef]
71. Aas, K.; Czado, C.; Frigessi, A.; Bakken, H. Pair-Copula Constructions of Multiple Dependence. *Insur. Math. Econ.* **2009**, *44*, 182–198. [CrossRef]
72. Steinschneider, S.; Brown, C. A Semiparametric Multivariate, Multisite Weather Generator with Low-Frequency Variability for Use in Climate Risk Assessments. *Water Resour. Res.* **2013**, *49*, 7205–7220. [CrossRef]
73. Alluri, A.; Aeidapu, M. Data-Driven Probabilistic Evaluation of Voltage Stability Limit Considering Joint Wind Uncertainties. *Int. J. Emerg. Electr. Power Syst.* **2026**. [CrossRef]
74. Lazović, Đ.; Đurišić, Ž. Stochastic Scheduling of a Sector-Coupled Virtual Power Plant Integrating Wind Generation and Thermal Energy Storage. *Energy Convers. Manag. X* **2026**, *30*, 101638. [CrossRef]
75. Yang, S.; Chen, G.; Chen, S. Copula-Accelerated ADMM Consensus for Distributed Stochastic Microgrid Dispatch. In *Proceedings of the 2026 Conference on Innovation and Technologies of Energy Storage (CITES), Hong Kong, China, 23–25 July 2026*; IEEE: New York, NY, USA, 2026; pp. 93–97. [CrossRef]
76. Fan, C.; Xiao, F.; Zhao, Y. A Short-Term Building Cooling Load Prediction Method Using Deep Learning Algorithms. *Appl. Energy* **2017**, *195*, 222–233. [CrossRef]
77. Fan, C.; Wang, J.; Gang, W.; Li, S. Assessment of Deep Recurrent Neural Network-Based Strategies for Short-Term Building Energy Predictions. *Appl. Energy* **2019**, *236*, 700–710. [CrossRef]
78. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; Adaptive Computation and Machine Learning Series; MIT Press: Cambridge, MA, USA, 2016; ISBN 978-0-262-03561-3.
79. Li, A.; Xiao, F.; Fan, C.; Hu, M. Development of an ANN-Based Building Energy Model for Information-Poor Buildings Using Transfer Learning. *Build. Simul.* **2021**, *14*, 89–101. [CrossRef]
80. Zhou, Z.-H. *Machine Learning*; Springer: Singapore, 2021; ISBN 978-981-15-1966-6.
81. Kohonen, T. Self-Organized Formation of Topologically Correct Feature Maps. *Biol. Cybern.* **1982**, *43*, 59–69. [CrossRef]
82. Goodfellow, I.J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Nets. In *Proceedings of the Advances in Neural Information Processing Systems*; Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2014; Volume 27.
83. Chen, Y.; Wang, Y.; Kirschen, D.; Zhang, B. Model-Free Renewable Scenario Generation Using Generative Adversarial Networks. In *Proceedings of the 2019 IEEE Power & Energy Society General Meeting (PESGM), Atlanta, GA, USA, 4–8 August 2019*; IEEE: New York, NY, USA, 2020; p. 1.
84. Arjovsky, M.; Chintala, S.; Bottou, L. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*; Precup, D., Teh, Y.W., Eds.; PMLR: London, UK, 2017; Volume 70, pp. 214–223.
85. Yang, T.; Xiong, T.; Qin, S.; Zou, H.; Huang, A.; Lou, S. A WGAN-GP Extreme Wind Power Scenario Generation Method Based on Persistent Homology Feature Constraints. In *Proceedings of the 2026 International Conference on Electrical, Control and Information Technology (ECITech), Nanjing, China, 20–22 March 2026*; IEEE: New York, NY, USA, 2026; pp. 753–758.
86. Yoon, J.; Jarrett, D.; van der Schaar, M. Time-Series Generative Adversarial Networks. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*; Curran Associates Inc.: Red Hook, NY, USA, 2019; pp. 5508–5518.
87. Lin, Z.; Jain, A.; Wang, C.; Fanti, G.; Sekar, V. Using GANs for Sharing Networked Time Series Data: Challenges, Initial Promise, and Open Questions. In *Proceedings of the ACM Internet Measurement Conference*; Association for Computing Machinery: New York, NY, USA, 2020; pp. 464–483.
88. Hu, R.; Chen, X.; Du, F. High-Fidelity Synthetic Meteorological Year Generation Model Integrating Physical Priors and Dual Attention. *Discov. Appl. Sci.* **2026**, *8*, 1298. [CrossRef]

89. Tian, C.; Li, C.; Zhang, G.; Lv, Y. Data Driven Parallel Prediction of Building Energy Consumption Using Generative Adversarial Nets. *Energy Build.* **2019**, *186*, 230–243. [CrossRef]
90. Ma, Z.; Mei, G.; Piccialli, F. An Attention-Based Cycle-Consistent Generative Adversarial Network for IoT Data Generation and Its Application in Smart Energy Systems. *IEEE Trans. Ind. Inform.* **2023**, *19*, 6170–6181. [CrossRef]
91. Baasch, G.; Rousseau, G.; Evins, R. A Conditional Generative Adversarial Network for Energy Use in Multiple Buildings Using Scarce Data. *Energy AI* **2021**, *5*, 100087. [CrossRef]
92. Zhang, Y.; Yang, Y.; Zhang, Z.; Jiao, F.; Zhang, X.; Ukil, A. Physically Consistent Sequence Generative Adversarial Networks for Wind Power Scenario Generation. *Sustain. Energy Grids Netw.* **2026**, *47*, 102393. [CrossRef]
93. Yang, L.; Fu, X. A Review of Photovoltaic Uncertainty Modeling Based on Statistical Relational AI. *Energies* **2026**, *19*, 1509. [CrossRef]
94. Rezende, D.J.; Mohamed, S.; Wierstra, D. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. Available online: https://arxiv.org/abs/1401.4082v3 (accessed on 10 September 2026).
95. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. *arXiv* **2022**, arXiv:1312.6114.
96. Fan, C.; Chen, M.; Tang, R.; Wang, J. A Novel Deep Generative Modeling-Based Data Augmentation Strategy for Improving Short-Term Building Energy Predictions. *Build. Simul.* **2022**, *15*, 197–211. [CrossRef]
97. Brégère, M.; Bessa, R.J. Simulating Tariff Impact in Electrical Energy Consumption Profiles With Conditional Variational Autoencoders. *IEEE Access* **2020**, *8*, 131949–131966. [CrossRef]
98. Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. *arXiv* **2020**, arXiv:2006.11239.
99. Dhariwal, P.; Nichol, A. Diffusion Models Beat GANs on Image Synthesis. *arXiv* **2021**, arXiv:2105.05233.
100. Capel, E.H.; Dumas, J. Denoising Diffusion Probabilistic Models for Probabilistic Energy Forecasting. In *Proceedings of the 2023 IEEE Belgrade PowerTech, Belgrade, Serbia, 25–29 June 2023*; IEEE: New York, NY, USA, 2023; pp. 1–6.
101. Rasul, K.; Seward, C.; Schuster, I.; Vollgraf, R. Autoregressive Denoising Diffusion Models for Multivariate Probabilistic Time Series Forecasting. *arXiv* **2021**. Available online: https://proceedings.mlr.press/v139/rasul21a/rasul21a.pdf (accessed on 21 November 2026).
102. Tashiro, Y.; Song, J.; Song, Y.; Ermon, S. CSDI: Conditional Score-Based Diffusion Models for Probabilistic Time Series Imputation. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*; Curran Associates Inc.: Red Hook, NY, USA, 2021; pp. 24804–24816.
103. Goyal, M.; Mahmoud, Q.H. A Systematic Review of Synthetic Data Generation Techniques Using Generative AI. *Electronics* **2024**, *13*, 3509. [CrossRef]
104. Ndayipfukamiye, T.; Ding, J.; Sarwatt, D.S.; Philipo, A.G.; Ning, H. Adversarial Defense in Cybersecurity: A Systematic Review of GANs for Threat Detection and Mitigation. Available online: https://arxiv.org/abs/2509.20411v2 (accessed on 18 September 2026).
105. Zhao, Z.; Shi, Y.; Wu, S.; Yang, F.; Song, W.; Liu, N. Interpretation of Time-Series Deep Models: A Survey. *arXiv* **2023**, arXiv:2305.14582.
106. NetCDF Climate and Forecast (CF) Metadata Conventions. Available online: https://cfconventions.org/cf-conventions/v1.6.0/cf-conventions.html (accessed on 14 November 2026).
107. Villani, C. *Optimal Transport*; Grundlehren der mathematischen Wissenschaften; Springer: Berlin/Heidelberg, Germany, 2009; Volume 338, ISBN 978-3-540-71049-3.
108. Box, G.E.P.; Jenkins, G.M.; Reinsel, G.C.; Ljung, G.M. *Time Series Analysis: Forecasting and Control*, 5th ed.; Wiley: Hoboken, NJ, USA, 2015; ISBN 978-1-118-67502-1.
109. Wilks, D.S. *Statistical Methods in the Atmospheric Sciences*, 3rd ed.; Academic Press: Cambridge, MA, USA, 2011; Volume 100, ISBN 978-0-12-385022-5.
110. Trenberth, K.E. Some Effects of Finite Sample Size and Persistence on Meteorological Statistics. Part I: Autocorrelations. *Mon. Weather Rev.* **1984**, *112*, 2359–2368. [CrossRef]
111. von Storch, H.; Zwiers, F.W. *Statistical Analysis in Climate Research*; Cambridge University Press: Cambridge, UK, 1999.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.