Stability Data-Driven Solar Forecasting for Grid Scheduling in the Galápagos Islands

1st Kerly Mishell Vaca Vallejo
Independent
0000-0002-2650-6400

2nd Byron Ernesto Vaca Barahona
ESPOCH
0000-0002-3622-0668

3rd Braulio Paúl Balseca Dahua
ESPOCH
0000-0002-7111-7949

***Abstract*—This paper presents a stability-certified data-driven solar forecasting and robust grid scheduling framework for islanded power systems, with application to the Galapagos Islands. Daily solar irradiance data from the NASA/POWER database are modeled using Hankel time-delay embedding and Dynamic Mode Decomposition (DMD), enabling low-order linear representations of residual irradiance dynamics. The learned operator is certified through discrete-time Lyapunov conditions formulated as Linear Matrix Inequalities (LMIs), ensuring stable multi-step forecast propagation. Forecasted irradiance is mapped to photovoltaic (PV) availability using performance-ratio standards and incorporated into a robust receding-horizon scheduling problem under uncertainty. The control layer adopts a scenario-based robust Model Predictive Control (MPC) formulation to hedge against PV forecast errors while respecting operational constraints. Simulation results demonstrate feasible and conservative grid-import schedules, consistent with weak-grid operational requirements. The proposed pipeline bridges Koopman-based data-driven modeling, convex stability certification, and robust optimization, offering a scalable framework for renewable integration in isolated grids.**

***Index Terms*—solar irradiance forecasting, Hankel embedding, dynamic mode decomposition (DMD), reduced-order modeling, Lyapunov stability, linear matrix inequalities (LMIs), robust grid scheduling, Galápagos Islands**

I. Introduction

The increasing penetration of photovoltaic (PV) generation in weak and isolated grids introduces significant operational challenges due to the intrinsic variability and uncertainty of solar irradiance. In islanded systems such as the Galapagos Islands, where generation margins are limited and grid flexibility is constrained, accurate forecasting and robust scheduling strategies are essential for secure operation. Solar forecasting methods include statistical approaches and modern machine learning models [1], [2], yet these can lack stability guarantees and interpretability when embedded in multi-step predictive-control loops.

Koopman operator theory provides a principled framework for representing nonlinear dynamics through linear evolution in lifted spaces. Dynamic Mode Decomposition (DMD), a computable approximation of Koopman dynamics, has gained prominence in the modeling of complex systems [3]–[5]. Variants such as compressed-sensing DMD [6] and Hankel time-delay embeddings [7], [8] substantially enhance modal resolution and capture temporal structure in real-world signals. Their relevance to renewable energy applications has been further supported by Koopman-based model reduction works in energy systems [9] and by recent applications of DMD directly to solar irradiance forecasting [10].

When DMD-derived models are propagated over multi-day horizons, ensuring model stability becomes essential. Stability certificates based on discrete-time Lyapunov functions and Linear Matrix Inequalities (LMIs) [11], [12] allow checking whether the learned operator can be safely rolled out, yielding controlled prediction error growth. This convex machinery naturally complements data-driven residual models, particularly in the context of Hankel–DMD, which yields low-order linear updates well suited for Lyapunov-based stability analysis.

At the control layer, Model Predictive Control (MPC) provides a structured approach to operational scheduling under constraints [13]. Robust MPC formulations—including min-max control [14], scenario-based MPC [15], [16], and survey treatments of robustness [17]—enable systematic handling of uncertainties in PV availability. Contemporary developments in data-driven MPC, including set-theoretic [18] and resilient formulations [19], as well as Gaussian-process-based cautious MPC [20], demonstrate the growing role of data-driven models in power system operations under uncertainty.

The NASA/POWER database [21], [22], which offers trustworthy worldwide measurements and has been verified across a variety of climatic regions [23], provides the solar resource that powers this framework. These products provide a reliable foundation for studies on the integration of renewable energy sources in isolated grids and medium-term forecasting.

In this work, an integrated forecasting and scheduling approach is proposed that includes: (i) a residual model of daily irradiance using a Hankel DMD; (ii) a stability guarantee of the learned operator using a Lyapunov/LMI approach; and (iii) a robust MPC formulation that takes into account constraints and uncertainty of PV power forecasts. The main contributions of this paper are:

- A stability-certified Hankel–DMD framework for multi-day irradiance forecasting.
- A separation-of-roles methodology where LMI-based stability ensures safe model propagation, while uncertainty is handled at the MPC layer through scenario-based robust optimization.
- The Galapagos Islands' end-to-end validation using NASA/POWER data shows cautious but workable grid-import scheduling that complies with weak-grid operational requirements..

II. Data and Preprocessing (Galápagos)

The NASA/POWER database provides daily solar resource measurements for a specific location in the Galapagos Islands. At daily resolution, the record spans January 2020 through December 2025. A daily proxy of the solar energy available to a photovoltaic plant, the modeling signal is the all-sky surface shortwave downwelling irradiance, `ALLSKY_SFC_SW_DWN` (kWh/m^2/day). The series is bounded and shows daily variability around a relatively stable level, as seen in Fig. 1, which is suitable for a low-order residual dynamics model.

A "calendar split" is used to prevent "look ahead": model identification is carried out on data prior to 2025-01-01, while forecasting and scheduling tests are carried out on days after 2025-01-01.

All model identification procedures are carried out on data up to 2025-01-01, and all reported forecasting and scheduling tests are conducted on days that follow in order to prevent look-ahead.

A. Cleaning and Missing Values

The flag value -999 may be used by NASA/POWER to encode missing entries. These samples are regarded as missing data, and `NaN` is used in their place. Time-based interpolation is used to fill in interior gaps in order to maintain the daily continuity that the delay-embedding construction requires. Forward/backward filling is used to prevent dropping entire days if there are missing values at the beginning or end of the record. Following cleaning, the data is converted to numeric format and sorted by date.

B. Scaling for Identification

To improve numerical conditioning in the Hankel–DMD identification, the irradiance series is standardized using training statistics only. Let y_k denote the cleaned irradiance at day k in physical units. Using the training subset, the mean μ and standard deviation σ are computed and the normalized signal is

$$\tilde{y}_k = \frac{y_k - \mu}{\sigma}. \tag{1}$$

All forecasts are mapped back to physical units using the inverse transform, which keeps the test period fully unseen during preprocessing.

C. Interface to Scheduling

In accordance with the data resolution, the control layer functions on a daily horizon. A straightforward capacity and performance-ratio mapping is used to translate irradiance forecasts into an average daily PV power (Section VI). A daily load signal $P_{\mathrm{load}}(k)$ is also necessary for grid import/export results; in the absence of measured demand, a proxy profile can be used for algorithmic validation; however, operational conclusions should be based on measured load.

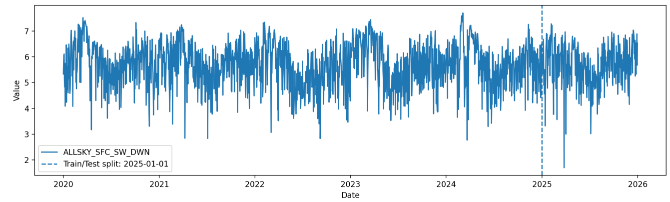

Figure 1: Daily irradiance series (`ALLSKY_SFC_SW_DWN`) and the calendar train/test split.

III. Data-Driven Forecasting Method (Hankel–DMD)

Daily irradiance is forecast using a low-order linear model identified directly from data through Hankel time-delay embedding and Dynamic Mode Decomposition (DMD). The main idea is to represent the recent history of the signal as an augmented state and learn a linear evolution rule in that embedded space. This provides a simple multi-step predictor that can be propagated over several days.

IV. Notation and Setup

This section fixes notation for the identification and forecasting stages and states the reduced-order residual dynamics used later in the stability and MPC layers.

A. Standardized signal and residual form

Let y_k denote the cleaned daily irradiance measurement at day k (Section II). Preprocessing produces the standardized series $\tilde{y}_k$ using the training mean μ and standard deviation σ as defined in (1). We use the same transformation for all channels and apply the inverse map to return forecasts to physical units.

To separate slow variations from short-term dynamics, a trend $\bar{y}_k$ is estimated over the training period and removed:

$$r_k = \tilde{y}_k - \bar{y}_k. \tag{2}$$

Then, the identification of the forescating model is r_k.

B. Delay embedding and Hankel construction

Given an embedding length L, we form the Hankel matrix from the residual series,

$$H = \begin{bmatrix} r_1 & r_2 & \cdots & r_{N-L+1} \\ r_2 & r_3 & \cdots & r_{N-L+2} \\ \vdots & \vdots & \ddots & \vdots \\ r_L & r_{L+1} & \cdots & r_N \end{bmatrix}, \tag{3}$$

and compute a rank-r approximation via truncated SVD,

$$H \approx U_r \Sigma_r V_r^{\top}. \tag{4}$$

As in standard Hankel–DMD, two shifted blocks are defined by

$$H_1 = H(:, 1 : end - 1), \qquad H_2 = H(:, 2 : end). \tag{5}$$

C. Reduced-order DMD operator and forecast recursion

The least-squares DMD operator that advances the embedded state is

$$A = H_2 H_1^{\dagger}, \tag{6}$$

where $(\cdot)^{\dagger}$ denotes the Moore–Penrose pseudoinverse. A reduced-order realization is obtained by projecting onto the dominant subspace,

$$\tilde{x}_{k+1} = \tilde{A}\tilde{x}_k, \qquad \tilde{A} = U_r^{\top} A U_r, \tag{7}$$

with $\tilde{x}_k \in \mathbb{R}^r$. Multi-step residual forecasts are produced by iterating (7) over the horizon and reconstructing r_{k+h} from the embedded state. Finally, the predicted irradiance in physical units is obtained by adding back the trend and inverting the standardization defined in (1).

D. Simulation protocol

All identification steps, i.e., estimation of trends, scaling, construction of Hankel matrices, and fitting of operators, operate only on the training set specified in Section II. Forecasting and robust MPC evaluation occur during the test period.

E. Hankel Time-Delay Embedding

A delay-embedded state is formed by stacking the last L standardized observations,

$$x_k = \begin{bmatrix} \tilde{y}_{k-L+1} \\ \tilde{y}_{k-L+2} \\ \vdots \\ \tilde{y}_k \end{bmatrix} \in \mathbb{R}^{pL}, \qquad k = L, \ldots, T. \tag{8}$$

This construction, which is frequently used to capture temporal structure through time delays, is equivalent to creating a block-Hankel matrix from the time series.

Using (8), snapshot matrices are defined as

$$\begin{aligned} X &= \begin{bmatrix} x_L & x_{L+1} & \cdots & x_{T-1} \end{bmatrix}, \\ X^+ &= \begin{bmatrix} x_{L+1} & x_{L+2} & \cdots & x_T \end{bmatrix}. \end{aligned} \tag{9}$$

with $X, X^+ \in \mathbb{R}^{pL \times N}$ and $N = T - L$.

The embedded state is assumed to follow an approximately linear evolution,

$$x_{k+1} \approx A x_k, \tag{10}$$

where the matrix $A \in \mathbb{R}^{pL \times pL}$ is estimated from data.

F. Reduced-Order DMD Identification

To obtain a compact model and reduce sensitivity to noise, a truncated SVD of X is used:

$$X = U \Sigma V^{\top}, \qquad X \approx U_r \Sigma_r V_r^{\top}, \tag{11}$$

where $U_r \in \mathbb{R}^{pL \times r}$ contains the leading r left singular vectors, $\Sigma_r \in \mathbb{R}^{r \times r}$ contains the corresponding singular values, and $V_r \in \mathbb{R}^{N \times r}$ contains the leading right singular vectors.

The reduced DMD operator is then computed as

$$\tilde{A} = U_r^{\top} X^+ V_r \Sigma_r^{-1} \in \mathbb{R}^{r \times r}. \tag{12}$$

When needed for rollouts in the full embedded coordinates, an associated lifted operator can be written as

$$A \approx U_r \tilde{A} U_r^{\top}. \tag{13}$$

The rank r controls the model complexity. In practice, L and r are selected based on validation performance and numerical stability.

G. Multi-Step Forecasting

Given the most recent L standardized observations at time k, the initial embedded state x_k is formed using (8). Multi-step forecasts are produced by iterating the learned dynamics:

$$\hat{x}_{k+h} = A^h x_k, \qquad h = 1, \ldots, H, \tag{14}$$

where H is the forecast horizon in days. The predicted measurement $\hat{\tilde{y}}_{k+h}$ is taken from the last block of $\hat{x}_{k+h}$ (for $p = 1$, the last entry). Predictions in physical units are obtained by inverting the standardization:

$$\hat{y}_{k+h} = \sigma\, \hat{\tilde{y}}_{k+h} + \mu. \tag{15}$$

H. Baselines and Accuracy Metrics

Forecast performance is evaluated on the test period using rolling predictions. Two simple baselines are used for reference:

- *Persistence:* $\hat{y}_{k+h} = y_k$ for all h.
- *Simple Moving Average (SMA):* $\hat{y}_{k+h} = \frac{1}{W} \sum_{j=0}^{W-1} y_{k-j}$ with window length W.

For a set of M evaluated points, the mean absolute error (MAE) and root mean squared error (RMSE) are computed as

$$\mathrm{MAE} = \frac{1}{M} \sum_{i=1}^{M} |y_i - \hat{y}_i|, \qquad \mathrm{RMSE} = \sqrt{\frac{1}{M} \sum_{i=1}^{M} (y_i - \hat{y}_i)^2}. \tag{16}$$

These metrics are reported across multiple horizons to characterize short- and multi-day forecast behavior.

V. Stability Assessment and Robust Certification

The learned residual dynamics must be rolled forward over a finite horizon for multi-step forecasting and day-ahead scheduling. Before using the identified residual operator inside MPC, we verify its stability because even tiny identification errors can add up over time. While the learned operator stays fixed, PV-availability scenarios (Section VI) handle the MPC robustness itself.

A. Nominal stability checks

The reduced-order residual model identified by Hankel–DMD:

$$\tilde{x}_{k+1} = \tilde{A}\tilde{x}_k, \tag{17}$$

with $\tilde{x}_k \in \mathbb{R}^r$. As a first check, we verify whether the eigenvalues of $\tilde{A}$ lie inside the unit disk, i.e., $\rho(\tilde{A}) < 1$. For the fixed finite-dimensional discrete-time linear model in (17), this spectral condition is theoretically equivalent to the existence of a quadratic Lyapunov certificate.

To obtain a constructive certificate, the quadratic Lyapunov condition by search a symmetric matrix $P \succ 0$:

$$\tilde{A}^\top P \tilde{A} - P \prec 0. \tag{18}$$

Feasibility of (18) certifies exponential stability of the identified residual dynamics and prevents open-loop rollouts from drifting when forecasts are propagated over multiple steps.

B. Robustness in scheduling vs. model certification

The robust component of the proposed pipeline is implemented at the *optimization* level rather than by modifying the learned dynamics. In particular, PV uncertainty is represented by a finite set of availability scenarios over the MPC horizon (Section VI). The controller then optimizes a worst-case objective across these scenarios while enforcing the grid and power-balance constraints for the selected robustness mode.

This design does not create a family $\{\tilde{A}_i\}$, and the DMD residual operator $\tilde{A}$ is *not* re-identified for each scenario. Instead, (18) is used to certify stability once for the fixed operator. This separation preserves the distinction in purpose: the scenarios account for PV forecasting errors at the scheduling stage, whereas the Lyapunov LMI certifies the multi-step predictability of the data-driven residual model.

C. Implementation details

The LMI in (18) is solved as a feasibility problem with decision variable $P = P^\top$ under the constraint $P \succ 0$. In the reported experiments, the condition is checked on the reduced operator $\tilde{A}$ associated with the chosen (L, r) pair used throughout the test period.

VI. Robust Grid Scheduling Under PV Forecast Uncertainty

The scheduling layer that converts the irradiance forecasts into a workable grid import/export plan is explained in this section. The only controllable variables are grid exchange and, when allowed, PV curtailment because no storage device is taken into account. Respecting grid limits while protecting against forecast uncertainty in PV availability is the aim.

A. PV Availability from Irradiance Forecasts

Let $\widehat{S}_{k+h}$ denote the h-step-ahead forecast of daily irradiance in kWh/m^2/day (e.g., `ALLSKY_SFC_SW_DWN`). A simple daily mapping is used to obtain the average available PV power:

$$\widehat{P}^{\mathrm{pv,avail}}_{k+h} = \frac{P_{\mathrm{dc}}\,\mathrm{PR}}{24}\,\widehat{S}_{k+h}, \tag{19}$$

where P_{dc} is the PV installed capacity (kWp) and PR is the performance ratio. The factor $1/24$ converts daily energy to average daily power. This approximation is suitable for daily scheduling and can be refined if sub-daily data are available.

B. Power Balance and Decision Variables

At each day k, the net grid exchange p^{g}_k (kW) is defined as positive for import and negative for export. The PV power used is denoted by p^{pv}_k (kW), with the option to curtail PV when needed to satisfy constraints. The daily load is P^{load}_k (kW). The balance constraint is

$$p^{\mathrm{g}}_k + p^{\mathrm{pv}}_k = P^{\mathrm{load}}_k. \tag{20}$$

For implementation convenience, the grid exchange can be split into import and export components,

$$p^{\mathrm{g}}_k = p^{\mathrm{imp}}_k - p^{\mathrm{exp}}_k, \qquad p^{\mathrm{imp}}_k \ge 0, \quad p^{\mathrm{exp}}_k \ge 0. \tag{21}$$

C. Operational Constraints

Grid exchange is limited by import and export capacities:

$$0 \le p^{\mathrm{imp}}_k \le P^{\mathrm{imp}}_{\max}, \qquad 0 \le p^{\mathrm{exp}}_k \le P^{\mathrm{exp}}_{\max}. \tag{22}$$

PV usage is constrained by the available PV:

$$0 \le p^{\mathrm{pv}}_k \le P^{\mathrm{pv,avail}}_k. \tag{23}$$

If export is not allowed, $P^{\mathrm{exp}}_{\max}$ is set to zero.

D. Uncertainty Model and Robust Constraints

PV availability is uncertain due to forecast errors. An uncertainty set is defined around the PV forecast:

$$P^{\mathrm{pv,avail}}_{k+h} \in \left[\widehat{P}^{\mathrm{pv,avail}}_{k+h} - \Delta_h,\ \widehat{P}^{\mathrm{pv,avail}}_{k+h} + \Delta_h\right], \tag{24}$$

where $\Delta_h \ge 0$ depends on the forecast horizon h. In practice, Δ_h can be estimated from out-of-sample residuals of the forecasting model (Section III).

A tractable robust formulation is obtained using a scenario approximation. Let $\{P^{\mathrm{pv,avail},(s)}_{k+h}\}_{s=1}^{N_s}$ denote a set of PV scenarios. The PV constraint is then enforced for all scenarios:

$$\begin{gathered} 0 \le p^{\mathrm{pv}}_{k+h} \le P^{\mathrm{pv,avail},(s)}_{k+h}, \\ \forall s \in \{1,\dots,N_s\},\ \forall h \in \{0,\dots,H-1\}. \end{gathered} \tag{25}$$

A conservative choice is the worst-case lower bound,

$$P^{\mathrm{pv,avail},(s)}_{k+h} = \max\{0, \widehat{P}^{\mathrm{pv,avail}}_{k+h} - \Delta_h\}, \tag{26}$$

which guarantees feasibility when realized PV is below the forecast within the assumed uncertainty band.

E. Receding-Horizon Optimization

At each day k, a finite-horizon optimization of length H is solved. The decision variables are $\{p^{\mathrm{imp}}_{k+h}, p^{\mathrm{exp}}_{k+h}, p^{\mathrm{pv}}_{k+h}\}_{h=0}^{H-1}$. A convex quadratic objective is used to reduce grid imports and avoid aggressive changes in grid exchange:

$$\begin{aligned} \min \ & \sum_{h=0}^{H-1} \left(w_g \left(p^{\mathrm{imp}}_{k+h}\right)^2 + w_e \left(p^{\mathrm{exp}}_{k+h}\right)^2 + w_c\, c_{k+h} \right) \\ & + w_{\Delta g} \sum_{h=0}^{H-2} \left(p^{\mathrm{g}}_{k+h+1} - p^{\mathrm{g}}_{k+h}\right)^2 \end{aligned} \tag{27}$$

$$\begin{aligned} \text{s.t.} \quad & p^{\mathrm{imp}}_{k+h} - p^{\mathrm{exp}}_{k+h} + p^{\mathrm{pv}}_{k+h} = P^{\mathrm{load}}_{k+h}, \quad h = 0,\dots,H-1, \\ & (22), \quad h = 0,\dots,H-1, \\ & (25), \quad h = 0,\dots,H-1. \end{aligned}$$

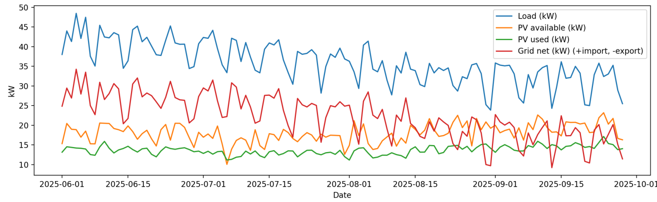

Figure 2: Daily power balance for the grid-only scheduling case: load, PV available, PV used, and net grid exchange.

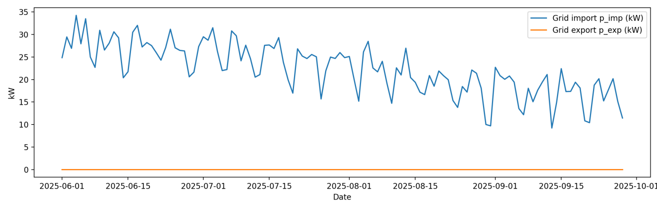

Figure 3: Grid import and export trajectories.

where $p^{\mathrm{g}}_{k+h} = p^{\mathrm{imp}}_{k+h} - p^{\mathrm{exp}}_{k+h}$, and $c_{k+h} \geq 0$ is a curtailment proxy that penalizes using less PV than forecasted:

$$c_{k+h} \geq \widehat{P}^{\mathrm{pv,avail}}_{k+h} - p^{\mathrm{pv}}_{k+h}, \qquad c_{k+h} \geq 0. \tag{28}$$

The weights $w_g, w_e, w_c, w_{\Delta g} \geq 0$ tune the trade-off between import reduction, export smoothing, curtailment preference, and ramping.

After solving (27), only the first-step decision is applied and the horizon is shifted forward, producing a receding-horizon policy.

F. Output Metrics and Plots

The net grid exchange p^{g}_k, as well as the import and export components, and the PV utilization, is used to summarize the scheduling behavior. Plots are provided for the following quantities: (i) daily power balance, (ii) import/export trajectory for the grid, and (iii) operational proxies such as the squared import and the grid ramp magnitude. The plots are provided in Figures 2, 3, and 4.

VII. Results and Discussion

This section reports the end-to-end experiment that links the Hankel–DMD irradiance forecast to robust grid-only scheduling for the Galápagos point case. A calendar split at 2025-01-01 is used, with the DMD model identified using $L = 60$ and $r = 25$. Grid-only MPC is simulated from 2025-06-01 for $N_{\mathrm{sim}} = 120$ days with horizon $H = 14$.

A. Forecast-to-PV mapping and availability

The daily irradiance forecast is mapped to an average available PV power using (19) with $P_{\mathrm{dc}} = 100$ kWp and $\mathrm{PR} = 0.8$. Figure 2 highlights two practical effects: (i) the forecast-driven PV profile captures the day-to-day variability that drives the schedule, and (ii) the PV *used* by the optimizer is sometimes below the mapped availability due to robust hedging and the curtailment proxy. The main averages for PV availability and PV usage over the test simulation are summarized in Table I.

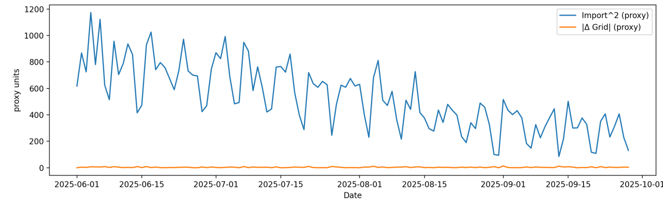

Figure 4: Operational proxies: squared grid import and absolute grid ramp magnitude.

B. Grid exchange under import/export limits

Figures 2 and 3 show that the schedule relies on grid imports throughout the simulated period, with no sustained surplus leading to export. This is consistent with the relative scale of PV and demand in the test configuration and with the conservative nature of the robust formulation, which discourages decisions that would be feasible only under optimistic PV realizations. Average and worst-case grid exchange values are reported in Table I.

C. Effect of hedging

Given the representation of PV uncertainty as a percentage band around the forecast, $\sigma_{\mathrm{pv}} = 0.15, \widehat{P}^{\mathrm{pv,avail}}$, and the safety factor $z = 1.64$, robustness is imposed by adopting a worst-case policy over $K = 7$ scenarios. In terms of its operation, this reduces the PV available to the optimizer to safely rely on over time. Thus, the controller takes a conservative approach by prioritizing import plans which are feasible rather than those which will exceed limits if PV generation underperforms.

D. Operational proxies and smoothness

The two proxies used to interpret the schedules are the magnitude of daily net grid ramps and the squared grid import. These two proxies are as presented in Figure 4. When the weak grid scenarios are considered, where large ramps are penalized operationally even in the absence of contractual power limit constraints, the weight of the ramp term in the cost function, i.e., the weight of the term $w_{\Delta g}$, is set to 0.2 in order to limit the sudden changes in the daily grid exchange as presented in the schedule. The squared import proxy increases with days that have higher demand-PV gaps, as presented in Fig. 3

E. Limitations: load data availability

To determine the grid import/export values as well as the aforementioned proxies for operation, the scheduling layer requires a daily load profile given by $P_{\mathrm{load}}(k)$. In the present study, a synthetic (proxy) daily load profile is employed in the assessment of the feasibility with PV uncertainty. Note that while it restricts the interpretation of the results in terms of import/export values as well as cost values, it is useful in the validation of the forecasting-to-control pipeline. The same procedure should be employed in a site-validated study with actual $P_{\mathrm{load}}(k)$.

Grid export is identically zero because the objective provides no incentive to inject energy (zero export credit and a

Table I: Averages over the 120-day simulation.

Metric	Value
Mean PV available $\overline{P}^{\mathrm{pv,avail}}$ (kW)	18.023
Mean PV used $\overline{P}^{\mathrm{pv,used}}$ (kW)	13.658
Mean grid import $\overline{P}^{\mathrm{imp}}$ (kW)	22.337
Max grid import $P^{\mathrm{imp}}_{\mathrm{max,obs}}$ (kW)	34.252
Mean grid export $\overline{P}^{\mathrm{exp}}$ (kW)	0
Max grid export $P^{\mathrm{exp}}_{\mathrm{max,obs}}$ (kW)	0

penalty on p_k^{exp}), while PV curtailment is permitted and lightly penalized. Moreover, the worst-case robust PV cap constrains p_k^{pv} to a conservative lower bound, further reducing surplus PV and yielding an effective self-consumption operating mode.

VIII. Conclusion

This paper presented a practical pipeline that connects data-driven solar forecasting with stability checks and robust grid scheduling for a point location in the Galápagos Islands. Daily NASA/POWER irradiance data from 2020–2025 were modeled using Hankel–DMD to produce multi-day forecasts. To support reliable propagation over several days, the learned dynamics were evaluated using discrete-time Lyapunov conditions, and a finite-set LMI formulation was outlined for cases where multiple operators are identified from rolling windows.

The forecasts were then mapped to an average daily PV availability and used to solve a robust receding horizon problem with import/export schedule optimization, considering limits and hedging against PV uncertainty. In the reported simulation, the schedule was found to be feasible under conservative PV forecasts and resulted in consistent import trajectories without the need for export.

The approach should be extended to hourly or sub-hourly resolution, where cloud-driven intermittency is more pronounced and where nonlinear or stochastic forecasting models may provide a better description of extreme irradiance fluctuations.

References

[1] S. Sobri, S. Koohi-Kamali, and N. A. Rahim, "Solar photovoltaic generation forecasting methods: A review," *Energy*, vol. 156, pp. 459–497, 2018.

[2] K. J. e. a. Iheanetu, "Solar photovoltaic power forecasting: A review," *Sustainability*, vol. 14, no. 24, 2022.

[3] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, "On dynamic mode decomposition: Theory and applications," *arXiv preprint arXiv:1312.0041*, 2013. [Online]. Available: https://arxiv.org/abs/1312.0041

[4] P. J. Schmid, "Dynamic mode decomposition and its variants," *Annual Review of Fluid Mechanics*, vol. 54, pp. 225–254, 2022.

[5] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.

[6] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Compressed sensing and dynamic mode decomposition," *Journal of Computational Dynamics*, vol. 2, no. 2, pp. 165–191, 2016.

[7] H. e. a. Asada, "Exact parallelized dynamic mode decomposition with hankel time-delay embedding," *Archive of Applied Mechanics*, 2025.

[8] D. e. a. Yang, "Synchronized ambient data-based extraction of interarea modes using hankel block matrices," *International Journal of Electrical Power & Energy Systems*, vol. 127, 2021.

[9] S. Peitz and S. Klus, "Koopman operator-based model reduction for switched-system control of energy applications," *Journal of Nonlinear Science*, vol. 30, pp. 1–34, 2020.

[10] Y. Zheng, J. Wang, and Y. Zhang, "Solar irradiance forecasting using improved dynamic mode decomposition," *Energy Reports*, vol. 8, pp. 970–979, 2022.

[11] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. SIAM, 1994.

[12] J. G. VanAntwerp and R. D. Braatz, "A tutorial on linear and bilinear matrix inequalities," *Journal of Process Control*, vol. 10, no. 4, pp. 363–385, 2000.

[13] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[14] P. O. M. Scokaert and D. Q. Mayne, "Min-max feedback model predictive control for constrained linear systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1136–1142, 1998.

[15] A. Nawaz, S. Wang, Z. Shafiq, Y. Zhang, and J. Wang, "Distributed mpc-based energy scheduling for islanded multi-microgrids," *Applied Energy*, vol. 331, p. 120317, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0306261922014258

[16] Y. Lin, H. Zhang, S. Wang, and Z. Li, "A scenario-based stochastic mpc for power systems under renewable uncertainties," *International Journal of Electrical Power & Energy Systems*, vol. 158, p. 110046, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061524002461

[17] A. Bemporad and M. Morari, "Robust model predictive control: A survey," *Robustness in Identification and Control*, 1999.

[18] F. Giannini and D. Famularo, "A data-driven approach to set-theoretic model predictive control for nonlinear systems," *Information*, vol. 15, no. 5, p. 281, 2024.

[19] G. Franzè, F. Giannini, D. Famularo, and F. Tedesco, "Output-based resilient model predictive control of cyber–physical systems under cover attacks," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 10 324–10 329, 2023.

[20] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious model predictive control using gaussian process regression," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2020.

[21] "Nasa power data access viewer," https://power.larc.nasa.gov/, 2024.

[22] "Nasa open data: Prediction of worldwide energy resources (power)," https://data.nasa.gov/dataset/prediction-of-worldwide-energy-resources-power, 2024.

[23] H. K. e. a. Tayyeh, "Analysis of nasa power reanalysis products," *Journal of Hydrology*, 2023.