Towards a Digital Twin of the European Power System

FEDERICO SENTINERI, CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy
GABRIELE PADOVANI, University of Trento, Italy
PAOLA NASSISI, CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy
SANDRO FIORE, University of Trento, Italy

As computing demands have soared over the past several years, sustainability has become a core priority across all domains of computer and computational sciences. Carbon-aware computing—scheduling workloads based on grid carbon intensity—has, among others, emerged as a key strategy to reduce the environmental footprint of digital systems. However, existing platforms (ENTSO-E, Electricity Maps, WattTime) aggregate data at national or zonal level, lacking the spatial granularity needed for fine-grained optimization, and do not model how electricity physically propagates through transmission networks. Our contribution is two-fold. We present a physics-informed digital twin of the European power system, integrating plant-level generation reconstruction, sector-decomposed demand modeling, and nodal power-flow simulation. We then derive H-REMD (High-Resolution Energy Mix Dataset), a 25km resolution hourly raster dataset of electricity mix and carbon intensity. Validation against French RTE reference data confirms physically consistent behavior across all pipeline stages. H-REMD could enable more accurate carbon-aware applications than current coarse-grained platforms, and additionally support future scenario analysis for sustainable planning of energy and digital infrastructures across Europe.

CCS Concepts: • **Computing methodologies → Modeling and simulation**; • **Hardware → Power and energy**; • **Applied computing → Physical sciences and engineering**; • **Social and professional topics → Sustainability**.

Additional Key Words and Phrases: carbon-aware computing, digital twin, carbon intensity, electricity, sustainability

1 Introduction

Recent advances in artificial intelligence, cloud computing, and large-scale digital infrastructures have significantly increased the energy demand of modern computational systems. In this context, the rapid expansion of geographically distributed AI data centers has accelerated industry interest in sustainable computing and green infrastructure management. Several approaches aim to reduce environmental impact by shifting computations temporally and geographically toward regions with cleaner energy mixes, enabling carbon-aware scheduling [21, 32], while interest has grown in frameworks able to improve the resource-efficiency of any number of downstream tasks. Current platforms such as ENTSO-E [12], Electricity Maps [6], and WattTime [24], provide useful carbon-intensity estimates, but rely mainly on national or zonal aggregations. This limits spatial resolution and fails to capture sub-regional variability in electricity mixes. These limitations can lead to suboptimal decisions in distributed computing systems, where accurate carbon-intensity estimates are critical for workload placement. Neighboring regions within the same market zone may exhibit substantially different effective carbon intensities due to local generation and cross-border electricity flows, which coarse-grained models cannot capture. To address these issues, this paper presents a physics-informed digital twin of the European power system, integrating generation units, energy demand, and the high-voltage transmission network. From this model, we derive H-REMD (High-Resolution Energy Mix Dataset), a spatially explicit hourly dataset of electricity mix and carbon intensity at 25 km resolution. The digital twin reconstructs plant-level generation from ENTSO-E data using also meteorological inputs, while demand is spatially disaggregated via energy demand raster model. These components are coupled through a PyPSA-based nodal power-flow simulation, ensuring physically consistent electricity mixing and cross-regional propagation. France is used as the primary validation case, with the framework designed to scale to the European level. Unlike existing datasets, H-REMD explicitly models electricity flows and spatial redistribution of energy mixes across the grid, paving the way for fine-grained carbon-aware optimization and scenario analysis[1].

While traditional platforms operate as static data trackers, collecting and displaying aggregated statistics, this framework proposes a physics-informed digital twin of the European power system; which ingests real-world grid telemetry and constructs a high-resolution virtual replica obeying power-flow physics. The output is H-REMD, a 25 km hourly raster dataset of electricity mix and carbon intensity derived from physically consistent simulations, which can then be utilized by downstream physical systems. This article is structured as follows: Section 2 explores the literature taken into consideration for this work; Section 3 gives an overview of the methodology employed to implement the digital twin and the resulting dataset; Section 4 explains the validation procedure for the development phases; and finally Section 5 draws our conclusions and explores the future developments enabled by this work.

2 Background and Related Works

In recent years, the field of carbon aware computing has been thoroughly explored from an academic point of view, to the point that several companies now employ production-ready approaches to reduce their carbon footprint. Modern carbon-aware computing exploits temporal and spatial variations in grid carbon intensity to reduce the environmental impact of computational workloads. While these approaches have demonstrated their effectiveness in reducing emissions [21], subsequent investigation highlights that accuracy is strongly correlated with the spatial granularity of the carbon-intensity data [25] used as source, and that a coarser representations can limit optimization quality. Current platforms such as ENTSO-E, Electricity Maps [5], and WattTime [27] have become

Authors' Contact Information: Federico Sentineri, federico.sentiner@cmcc.it, CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy; Gabriele Padovani, gabriele.padovani@unitn.it, University of Trento, Italy; Paola Nassisi, paola.nassisi@cmcc.it, CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy; Sandro Fiore, sandro.fiore@unitn.it, University of Trento, Italy.

---

[1] Due to the validation process, currently performed on a subset of the data, both the codebase and the dataset created will be open sourced in a future release of the digital twin.

very popular by providing carbon-intensity estimates at national or zonal scale, using aggregated production data or flow-based approximations, but share three main limitations: (i) spatial aggregation that hides sub-regional variability, (ii) simplified treatment of cross-border electricity flows without full network physics, and (iii) limited support for future scenario analysis [25]. Tackling these issues from a different point of view, datasets such as the Global Power Plant Database (GPPD) [31] provide global plant-level information, but generation data is only partially available. To overcome this limitation, several studies estimate electricity generation using plant characteristics combined with meteorological reanalysis data, improving coverage but still facing uncertainties, and estimation error. Physics-based power-system models such as PyPSA [2] and PyPSA-Eur [13], on the other hand, track optimal power flow and allow for long-term energy analysis, although, they typically rely on simplified network representations to ensure computational tractability. On this trajectory, recent works have improved spatial detail by reconstructing transmission networks from OpenStreetMap data within PyPSA-Eur workflows [30]. These models achieve higher geographic fidelity while remaining consistent with ENTSO-E statistics, enabling more realistic representations of cross-border flows and infrastructure topology. Beyond generation data, the spatial distribution of electricity demand remains a significant challenge for high-resolution modeling. Traditional methods for disaggregating national load often rely on proxies such as population density or nightlight satellite imagery to estimate sub-regional consumption patterns [1]. While useful, these proxies fail to capture the distinct spatial footprints of industrial, residential, and transport sectors, necessary for improving socio economic simulations and other types of downstream applications. While current models either provide high-level historical accuracy or low-resolution physical forecasting [32], a unified framework that combines all the aforementioned aspects is still not available.

3 Digital Twin and H-REMD Dataset

This section describes the architecture of the digital twin and the methodology adopted to construct a high-resolution model of the European power system, from which H-REMD is derived.

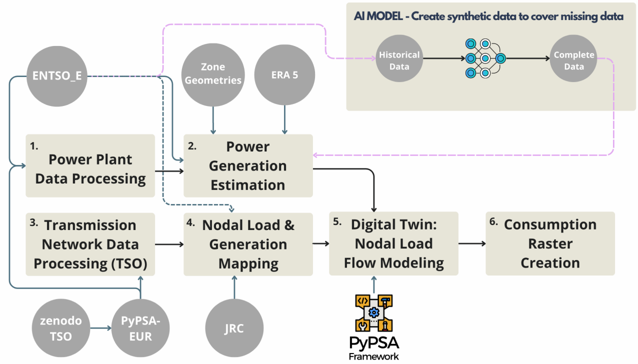

Fig. 1. High level overview of the pipeline's data sources, usage and flow through the six main stages.

Figure 1 displays an high level overview of the phases included in the dataset generation pipeline. The digital twin combines heterogeneous energy datasets, meteorological information, and physics-based power-flow simulation to compute the spatio-temporal evolution of the electricity mix and its associated carbon intensity. The overall process is composed of four main stages: (i) plant-level generation reconstruction, (ii) demand modeling, (iii) transmission network reconstruction, and (iv) power-flow simulation and H-REMD derivation.

3.1 Plant-Level Generation Matching

The first stage focuses on estimating hourly electricity generation for individual power plants. The reconstruction process integrates electricity production data from ENTSO-E and geolocalized plant metadata from the PyPSA-Eur ecosystem. A preliminary survey of available datasets — including the ENTSO-E Transparency Platform [7], the Joint Research Centre (JRC) Power Plants Database [15], the Global Power Plant Database [28], Open Power System Data [17], and the Global Energy Monitor [8] — revealed significant heterogeneity in plant counts, spatial accuracy, and capacity values. The PyPSA Power Plant Matching dataset [2, 9] was selected as the primary source due to its improved spatial and capacity consistency across Europe. ENTSO-E generation-unit entries were linked to geolocated plants in the PPM catalogue using direct code matching, nearest-neighbour geocoding via OpenStreetMap[18], and consistency checks on plant metadata. For renewable sources — wind, hydro and solar in particular — direct plant-level measurements are often unavailable. Generation is therefore estimated from ERA5 reanalysis data [10, 20], following established system-level renewable energy modeling approaches. For photovoltaic generation, irradiance is projected onto the plane-of-array using classical solar geometry [16] and converted into electricity output using a PVWatts-type performance model [4, 23]. Wind output is derived from ERA5,using a power-law extrapolation and a standard turbine power-curve formulation [3]. To preserve consistency with national totals, residual discrepancies between reconstructed (G) and observed zonal generation (G^{obs}) are redistributed proportionally to plant capacity:

$$\Delta G_{z,t} = G_{z,t}^{\text{obs}} - \sum_{i \in z} G_{i,t} \tag{1}$$

where t indicates the current time step and z represents the selected zone, ensuring exact aggregation consistency at each hour. Reconstructed profiles are finally validated by comparison with regional production statistics from the French transmission system operator [22].

3.2 High-Resolution Demand Reconstruction

The second stage reconstructs electricity demand at high spatial resolution through a set of raster-based modeling techniques. Since detailed consumption measurements are generally unavailable at fine spatial granularity, the framework disaggregates national demand values onto a regular grid representation. For this, two complementary data sources are utilized: the hourly time series for the national demand are obtained from ENTSO-E and are treated as a hard constraints, while the spatial distribution is derived from the annual electricity raster provided by the JRC [11], which describes the pattern of total electricity consumption across Europe at 0.25° resolution. The disaggregation is performed starting from the ENTSO-E national total $D^{\text{obs}}_{z,t}$, which is distributed across the raster in proportion to the JRC annual values. This operation is repeated for each isolated bidding zones, as often electricity does not freely flow through these borders, and for every hour. Formally, for each cell $i \in z$, a static spatial weight is computed as showin in Eq. 2a, where R_i is the JRC annual demand value at cell i, so that the sum is equal to one.

$$w_i = \frac{R_i}{\sum\limits_{j \in z} R_j} \tag{2a}$$

$$D_{i,t} = D^{\text{obs}}_{z,t} \cdot w_i \tag{2b}$$

The hourly demand allocated to cell i at hour t is then calculated as shown in Eq. 2b. By construction, summing $D_{i,t}$ over all cells in zone z recovers $D^{\text{obs}}_{z,t}$ at every hour, while the spatial distribution of demand within the zone remains proportional to the JRC raster at all times. Zones for which the JRC raster provides insufficient coverage (e.g. fewer than a minimum number of valid pixels, or coverage fraction below a threshold) are excluded from the output. Once the rasterized demand is calculated, each cell is associated with the nearest transmission node of the electrical network. The resulting nodal demand is obtained by aggregating the demand of all raster cells assigned to the same node, and constitutes the input to the PyPSA power-flow simulation described in Section 3.4. As an alternative spatial proxy, the PyPSA-Eur demand estimate [13] was evaluated as a drop-in replacement in Eq. 2a. The reliability of the JRC raster as a spatial proxy is not uniform across regions: while most provinces show weights consistent with independent demand data, some are systematically over- or under-estimated. As this bias propagates into nodal demand and distorts the power-flow simulation, a corrective rescaling based on RTE provincial proportions was applied to the French raster; extending this correction to the rest of Europe, where independent ground truth is generally unavailable, is left for future work.

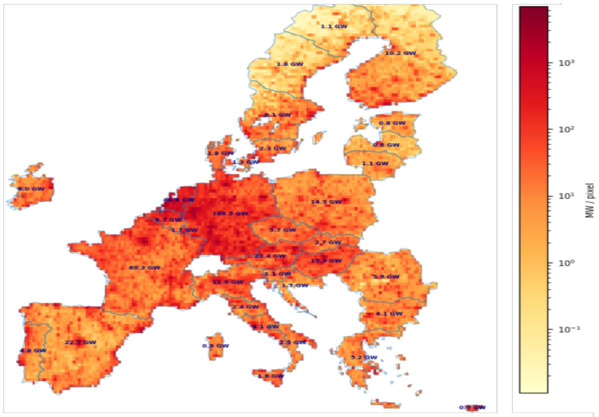

Fig. 2. Result of energy demand generated raster – 2019-01-01 00:00:00 CET. The figure also highlights clear spatial patterns, with areas of higher residential demand corresponding to major cities.

3.3 Transmission Network Reconstruction

The third stage consists in constructing a simplified representation of the European transmission network suitable for digital twin simulation. The network topology is derived primarily from PyPSA-Eur and additional transmission infrastructure datasets obtained from prior transmission network reconstruction studies [29].

The electrical network is modeled through the interplay between three main components: generators, which represent electricity production units; transmission lines (AC) and links (DC), consisting in high-voltage electricity transport infrastructure; and buses, which identify network nodes where electricity flows converge and mix, determining how electricity generated from different sources propagates through the system. In order to associate the newly proposed power plants to existing busses, the generators were inspected, finding that over 95% of power plants are connected to the geographically nearest bus. Justified by this insight, all remaining power plants are connected to the respective closest bus. The full scale network of buses and transmission lines is then simplified through a clustering procedure, aggregating neighboring nodes into representative macro-nodes. This operation reduces the computational complexity of the digital twin while preserving the large-scale structure of the network. Generation units are then reassigned to the corresponding clustered nodes according to their spatial location, enabling scalable power-flow simulations while maintaining a physically meaningful approximation of electricity propagation across the network.

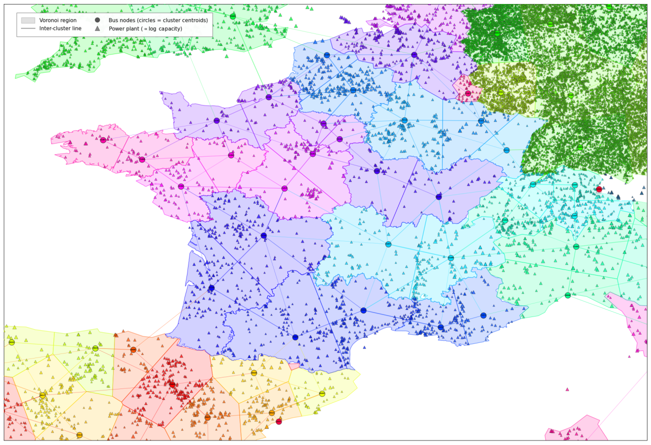

Fig. 3. Cluster area delineation with associated power plants and transmission network. The cluster boundaries closely follow provincial and national borders, while plant-to-cluster association depends purely on geographic location. Boundaries between clusters belonging to the same province follow a linear Voronoi-based separator.

3.4 Power-Flow Simulation and H-REMD Derivation

The final stage consists in instantiating the digital twin by integrating reconstructed generation, nodal demand, and transmission network information into the PyPSA simulation framework.

For each time step, the digital twin receives: (i) the hourly production profiles of individual generation units (Section 3.1); (ii) the nodal demand profiles (Section 3.2); and (iii) the clustered transmission network (Section 3.3). Based on these inputs, the framework computes power flows across the network, while accounting for the topology and physical constraints. Electricity generated by power plants propagates according to nodal demand, calculating imports, exports, and the mix across interconnected regions. The energy composition is therefore not determined solely by local generation, but also by imported electricity originating from neighboring regions.

Power flows are computed using a linear formulation, solving Kirchhoff's laws given the reconstructed generation and demand as fixed injections. This framework aims to reproduce how electricity physically flows given observed conditions, rather than to prescribe an optimal dispatch under a cost or emissions objective. An optimization-based dispatch was evaluated but discarded: incomplete sub-threshold generation coverage (Section 4.2) and spatial demand uncertainty (Section 4.3) leave several nodes underconstrained, causing the optimizer to redistribute the missing balance through extreme marginal adjustments rather than a feasible dispatch. This is a temporary solution; future work will combine optimized and linear power flow once generation and demand coverage are sufficiently complete to keep the optimization well-constrained. Parameters are not available on a per-line level and are approximated by standard per-voltage values, consistent with the line-type mapping approach adopted in PyPSA-Eur [13, 29]. DC links are excluded from the current network, as their transport-style flow formulation is incompatible with the impedance-based solver used for AC lines. Excluded connections either terminate in countries outside the modeled scope (e.g., UK, Nordic countries), while others are geographically distant from the French validation area, with no measurable effect on inter-regional flows (Section 4.4).

The per-bus energy mix is then spatially interpolated onto a regular raster grid via exponentially weighted Soft Voronoi [19] procedure, producing the final dataset, with hourly temporal resolution and 25 km spatial resolution, currently covering France and designed to extend to the full European network. Though the digital twin does not directly alter the state of the European transmission network, it acts as an intermediary controller. It falls into the category of Dynamic Digital Twin [14], converting raw national telemetry into a 25 km sub-regional raster. This higher resolution provides physically accurate ground truth data for downstream execution, and ensures that the virtual model mirrors the characteristics of the real infrastructure [26].

4 Data Validation

The validation of the H-REMD framework is organized under three distinct topics: (i) generation, (ii) demand, and (iii) power-flow within the digital twin. Each section focuses only on the provincial level for France, which was selected as the validation case due to the availability and granularity of RTE provincial-level data, combined with its diversified generation mix and central role in cross-border electricity exchange across Europe. Other well-instrumented countries, such as Germany, represent natural candidates for future validation efforts.

4.1 Impact of Clustering on Performance

As mentioned, to ensure computational feasibility for large-scale simulations, the digital twin utilizes a clustering procedure to aggregate buses geographically. The sensitivity of the model to the number of clusters was evaluated by measuring execution time against two entropy measurements. Shannon and Local Entropy were chosen to quantify respectively the diversity and spatial heterogeneity of the energy mix captured at different clustering levels.

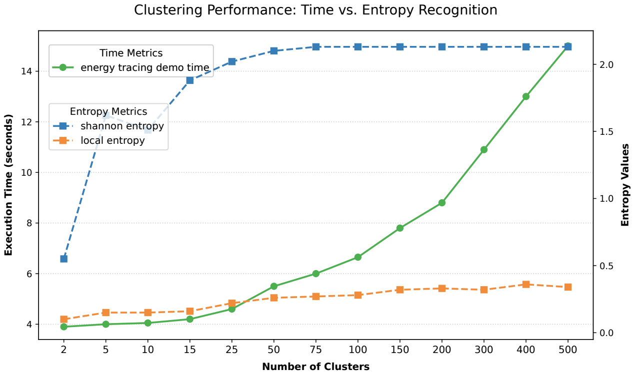

Fig. 4. Execution time of energy-tracing simulation (green) alongside Shannon entropy (blue) and local entropy (orange) for simulation runs with increasing number of clusters.

Based on the results shown in figure 4, the range of 100 to 200 clusters was identified as the "elbow" region. This range provides a high degree of entropy recognition while avoiding the exponential increase in execution time observed at higher node counts. A more thorough evaluation of the clustering method will be conducted with a set of downstream use cases, but to obtain a preliminary assessment of the information present in the raster, this approach was considered satisfying.

4.2 Generation Validation

The generation validation correlates plant-level estimates against RTE provincial data for each energy source and province. Nuclear power reconstruction exhibits the highest accuracy, due to the limited number of facilities and their monitoring in ENTSO-E. Wind, hydro, and solar show strong temporal correlation but higher error percentage compared to nuclear. However, since their absolute generation volumes are significantly lower, these errors have minimal impact on the total energy mix. Moreover, the generation is constrained by national-level totals, ensuring energy composition is preserved and errors only affect spatial allocation. Thermal generation presents the largest discrepancies. For large capacity plants (> 100 MW), reconstruction is reliable, as observed in provinces like Nor, PACA, and GE. However, many provinces exhibit a nearly constant thermal production in RTE not captured by the framework. This can be attributed to small-scale units (< 100 MW) in urban areas, not registered in ENTSO-E. Bre, shown in Figures 5 and 6, exemplifies this limitation: RTE reports steady thermal output with no corresponding large plant in the reconstruction, as

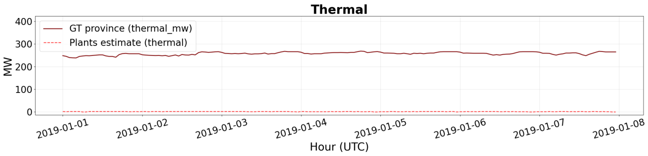

Fig. 5. Thermal generation in Bre: RTE ground truth vs. plant-level estimate. The absence of monitored thermal plants above the sub-threshold cutoff results in an estimated production of zero, despite RTE reporting a non-negligible baseline output in the province.

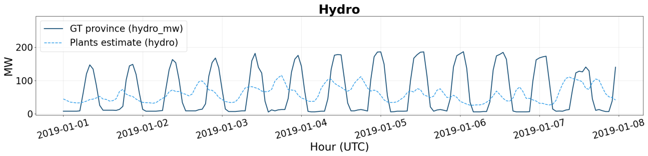

Fig. 6. Hydro generation in Bre: RTE ground truth vs. plant-level estimate. Small-scale hydro plants are influenced by the behavior of large national-level hydro facilities: since the estimation procedure starts from national totals dominated by major plants, their temporal pattern is imposed on small-scale plants as well.

production originates entirely from sub-threshold distributed co-generation units outside ENTSO-E monitoring; hydro generation in the same province, lacking direct measurements and instead estimated from national totals, is similarly more exposed to error when local volumes are small. Table 1a reports mean RMSE values on total generation of 13.2%, with thermal contributing the largest absolute error (117 MW) due to the unmonitored sub-threshold plants. The highest relative errors are observed in provinces with very low total generation (e.g., BFC, Bre, PDL), since small spatial allocation discrepancies translate into large percentage errors where production is already low. To prevent these provinces from dominating the aggregate metric, the reported total RMSE is weighted by each province's total generation, so errors in high-production provinces contribute proportionally more to the overall 13.2% figure.

4.3 Demand Validation

For validation of the electricity demand, the spatial disaggregation of hourly consumption was assesses across French provinces. More specifically, the raster-aggregated demand was compared against data provided by RTE at the provincial level for the full year 2019. For each province, the raster cumulate value at each hourly time step is computed by summing the demand at each cell of the province. The resulting provincial time series is then compared against the RTE consumption via Pearson correlation r, mean bias, and normalized root-mean-square error (NRMSE%). Table 1c, computed prior to the manual raster correction (Section 3.2), reveals substantial asymmetries: IdF is overestimated by +53.9% and PACA underestimated by −43.0%, both well above the 20.2% mean NRMSE, confirming non-uniform reliability of the JRC weights and motivating the correction. The results confirm that the disaggregation model reproduces the temporal dynamics of regional consumption with high fidelity across a majority of provinces. Pearson correlation coefficients are consistently high, which is expected, as the temporal load is primarily shaped by ENTSO-E national totals, preserved as a hard constraint (Eq. 2b).

Although the disaggregation methodology demonstrates robust performance in capturing temporal dynamics, the observed discrepancies in bias and NRMSE suggest that the source of the error cannot be uniquely attributed to the spatial weighting scheme. To investigate this issue, the hourly demand was additionally reconstructed from *PyPSA-Eur*, combining national energy totals, temporal profiles, spatial allocation factors, climate and socio-economic indicators. Validation against RTE provincial data, however, showed even larger errors. This indicates that inconsistencies may arise either from limitations in the RTE provincial data, or from structural inaccuracies in the spatial allocation (including, but not limited to, JRC-derived weighting factors). Consequently, the overall framework appears consistent at the temporal level, but spatial allocation uncertainties remain unresolved and warrant further investigation.

4.4 Power-Flow and H-REMD Validation

The validation of the power-flow simulation is performed by comparing inter-regional electricity exchanges simulated by the digital twin against reference measurements provided by RTE. For each pair of adjacent French provinces, all internal buses are aggregated into a single representative node, and the net inter-regional flow is computed. RTE provides a single aggregated import/export value per province, rather than bilateral flows between specific province pairs. Consequently, validation requires summing all simulated flows to and from neighboring provinces into a single net exchange value for comparison against the RTE reference. This constitutes a methodological limitation: a pairwise comparison of province-to-province flows would be preferable, but the corresponding ground-truth data is not available. Table 1b reports validation metrics for inter-regional exchanges across all 12 French provinces. Mean correlation is $r = 0.807$, with MAE ranging from 104 MW (BFC) to 4782 MW (GE). The high correlation confirms that the digital twin correctly models the temporal dynamics of cross-border flows, and that the power-flow simulation produces physically consistent inter-regional exchanges when supplied with accurate demand inputs.

As shown in Figure 7, the framework provides direct access to sub-regional carbon intensity at 25 km resolution, and neighboring areas within the same bidding zone can exhibit substantially different carbon intensities, a feature which is unavailable at coarser granularity.

The digital twin further enables scenario analysis by allowing direct manipulation of its components: modifying the demand raster or the network topology, and injecting synthetic generation units at specific buses, allows for simulations of how the energy mix and cross-border flows adapt to meet demand under hypothetical conditions.

5 Conclusions and Future Work

In this paper, we present an initial work towards a physics-informed digital twin of the European power system, integrating plant-level generation reconstruction, demand modeling, and nodal power-flow

Table 1. Comprehensive validation metrics for French provinces across generation, inter-regional exchange, and demand stages.

(a) RMSE (MW) by resource and % error on provincial total.

Prov.	**Hyd**	**Nuc**	**Sol**	**Ther**	**Wind**	**Tot**	**RMSE**
ARA	354	71	37	40	119	14476	4.3
BFC	46	—	13	51	25	349	38.6
Bre	70	—	22	256	51	416	96.1
CVL	17	29	13	76	63	10628	1.9
GE	355	3460	47	70	136	14586	28.1
HdF	2	48	14	191	114	7813	4.7
IdF	8	—	2	96	9	444	25.7
Nor	14	1063	4	51	24	7442	15.3
NAq	163	70	95	69	126	7436	8.8
Occ	254	6	57	91	216	4442	13.9
PACA	225	—	99	239	23	2401	24.4
PDL	4	—	34	168	73	678	40.9
Mean	**126**	**820**	**37**	**117**	**80**	**5909**	**13.2**

(b) Inter-regional exchange validation (r: Pearson index, P90: 90th percentile error).

Prov.	r	**MAE**	**RMSE**	**MAPE**	**Bias**	**P90**
ARA	0.687	948	1291	34.3	−787	2379
BFC	0.930	104	130	4.1	+61	220
Bre	0.881	219	271	8.1	+168	455
CVL	0.838	114	145	1.5	−52	244
GE	0.735	4782	4802	55.5	+4782	5240
HdF	0.675	856	1117	80.4	−583	1993
IdF	0.839	633	768	7.0	+11	1228
Nor	0.890	1147	1168	32.9	+1147	1410
NAq	0.741	748	846	1154.6	−725	1250
Occ	0.284	914	1067	44.4	−879	1629
PACA	0.585	972	1062	26.5	−917	1591
PDL	0.795	241	288	7.5	+81	452
Mean	**0.807**	**1057**	**1330**	**121.4**	+352	**1508**

(c) Per-province demand validation metrics (2019 data).

Prov.	r	**Bias%**	**NRMSE**
ARA	0.935	−0.7	7.6
BFC	0.927	+1.5	8.6
Bre	0.897	−4.8	12.3
CVL	0.920	−8.3	13.2
GE	0.914	−0.1	8.3
HdF	0.917	+1.1	8.8
IdF	0.911	+53.9	55.8
Nor	0.918	−6.7	11.0
NAq	0.914	−25.6	27.6
Occ	0.884	−16.3	19.9
PACA	0.831	−43.0	44.4
PDL	0.917	−21.2	24.4
Mean	**0.907**	−5.9	**20.2**

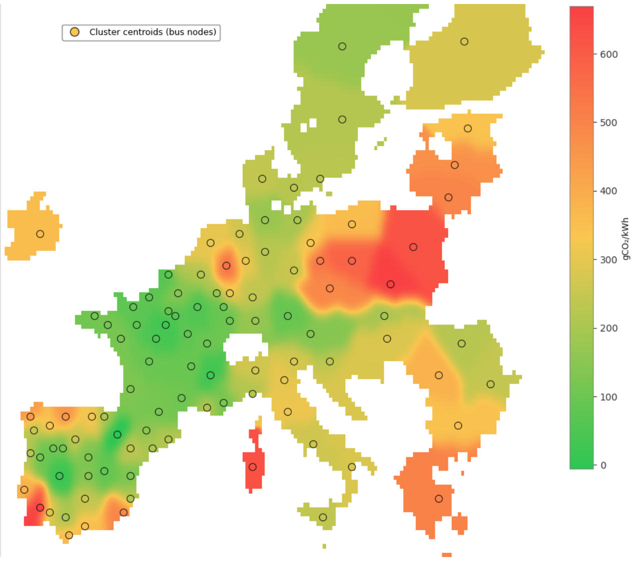

Fig. 7. H_REMD: CO_2 intensity raster derived from the digital twin via Soft Voronoi interpolation. Sub-regional variability within bidding zones is clearly visible, illustrating the limitations of national-level aggregations.

simulation. From this framework, H-REMD is derived: a 25 km resolution hourly raster of carbon intensity. Validation against French RTE data confirms physically consistent behavior across all pipeline stages, and unlike existing coarse-grained platforms, the digital twin explicitly models cross-border electricity flows, producing more accurate carbon-intensity information, to be used in downstream tasks such as workload scheduling and datacenter placement.

The proposed work is not without limitations. Downstream applications defined are currently limited to the framework's scenario-analysis capability, since the digital twin relies on a descriptive linear power-flow solver. Consequently its capabilities are limited to scenarios such as physical propagation of energy after the creation of hypothetical starting scenarios, or to the injection of a static generation profile for forecast generation.

Additionally, the current approach relies on the availability and accuracy of the JRC demand raster, so in regions with missing data—such as the Balkan states—the output raster cannot be computed. The framework is also limited to the European network, excluding extra-European import/export flows. Nevertheless, the highly structured format of this dataset lends itself to machine learning applications, suggesting that spatial data gaps could be resolved through predictive inference. Furthermore, a more comprehensive validation process for the entirety of the dataset has to be conducted, including all zones and benchmarking the performance of ML inference in areas in which data is not available.

Future work, moreover, will focus on the decomposition of demand into sector-level components, which will allow for socio-economic scenario analysis; and on the integration the distribution network into the clustering procedure, for improved definition of the supply zones associated with each bus.

Ultimately, H-REMD provides an important tool for the future architecture of next-generation European cloud infrastructures, offering high temporal and spatial resolution data for fine-grained carbon-aware workload scheduling and other sustainability-focused applications.

This work is considered a foundational phase: once the limitations discussed are addressed —namely the integration of an optimization-based solver alongside the linear power flow, improved generation and demand estimates— applied scenarios will be explored, such as evaluating the grid impact of new generation siting or alternative network configurations. Moreover, beyond carbon-aware scheduling, the same high-resolution mix and flow information could support related use cases such as siting analysis for new data centers, grid-aware demand response, and stress-testing of local infrastructure under increasing electrification of compute.

References

[1] Roberto Antonietti, Marco Bettiol, Eleonora Di Maria, and Luca Vargiu. 2025. *Data Centers' Sustainability in the EU.* GEOTWIN Policy Brief 1/2025. University of Padova, Department of Economics and Management "Marco Fanno".

[2] Tom Brown, Jonas Hörsch, and David Schlachtberger. 2018. PyPSA: Python for Power System Analysis. *Journal of Open Research Software* 6, 1 (2018), 4. doi:10.5334/jors.188

[3] Tony Burton, David Sharpe, Nick Jenkins, and Ervin Bossanyi. 2011. *Wind Energy Handbook* (2 ed.). Wiley.

[4] Aron P. Dobos. 2014. *PVWatts Version 5 Manual.* Technical Report NREL/TP-6A20-62641. National Renewable Energy Laboratory.

[5] Electricity Maps. 2025. Carbon Intensity API. https://www.electricitymaps.com. Accessed: 2025.

[6] Electricity Maps. 2026. Real-Time CO_2 Emissions of Electricity Consumption. https://app.electricitymaps.com/ Accessed: 2026.

[7] ENTSO-E. 2026. Transparency Platform. https://transparency.entsoe.eu/. Accessed: 2026.

[8] Global Energy Monitor. 2026. Global Power Plants Tracker. https://globalenergymonitor.org/. Accessed: 2026.

[9] Fabian Gotzens, Heidi Heinrichs, Jonas Hörsch, and Fabian Hofmann. 2019. Performing energy modelling exercises in a transparent way — the issue of data quality in power plant databases. https://github.com/PyPSA/powerplantmatching. *Energy Strategy Reviews* 23 (2019), 1–12.

[10] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* 146, 730 (2020), 1999–2049. doi:10.1002/qj.3803

[11] Ignacio Hidalgo Gonzalez, Hrvoje Medarac, and Andreas Uihlein. 2021. Data documentation for the Energy and Industry Geography Lab. (2021).

[12] Lion Hirth, Jonathan Mühlenpfordt, and Marisa Bulkeley. 2018. The ENTSO-E Transparency Platform–A review of Europe's most ambitious electricity data platform. *Applied energy* 225 (2018), 1054–1067.

[13] Jonas Hörsch, Fabian Hofmann, David Schlachtberger, and Tom Brown. 2018. PyPSA-Eur: An open optimisation model of the European transmission system. *Energy Strategy Reviews* 22 (2018), 207–215. doi:10.1016/j.esr.2018.08.012

[14] IBM. 2024. What is a digital twin maturity model? https://www.ibm.com/think/topics/digital-twin-maturity-model. Accessed: 2026-06-30.

[15] Joint Research Centre (JRC). 2026. JRC Open Power Plants Database (JRC-PPDB-OPEN). https://data.jrc.ec.europa.eu/dataset/9810feeb-f062-49cd-8e76-8d8cfd488a05. Accessed: 2026.

[16] B. Y. H. Liu and Richard C. Jordan. 1963. The interrelationship and characteristic distribution of direct, diffuse and total solar radiation and its application to solar energy estimation. *Solar Energy* 7, 2 (1963), 1–19. doi:10.1016/0038-092X(63)90062-6

[17] Open Power System Data. 2026. Conventional Power Plants. https://open-power-system-data.org/. Accessed: 2026.

[18] OpenStreetMap contributors. 2026. Planet dump retrieved from https://planet.openstreetmap.org. https://www.openstreetmap.org.

[19] Ravi Peters, Hugo Ledoux, and Martijn Meijers. 2014. A Voronoi-based approach to generating depth-contours for hydrographic charts. *Marine Geodesy* 37, 2 (2014), 145–166.

[20] Stefan Pfenninger and Iain Staffell. 2016. Long-term patterns of European PV output using reanalysis data. *Energy* 114 (2016), 1251–1265. doi:10.1016/j.energy.2016.08.060

[21] Ana Radovanović, Ross Koningstein, Ian Schneider, Binzhao Chen, Alexandre Duarte, Biniv Roy, and Walfredo Cirne. 2023. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems* 38, 2 (2023), 1270–1280. doi:10.1109/TPWRS.2022.3168153

[22] RTE. 2026. Regional electricity data – écoMix. https://www.rte-france.com/en/data-publications/eco2mix/regional-data. Accessed: 2026.

[23] Eva Skoplaki and John A. Palyvos. 2009. On the temperature dependence of photovoltaic module electrical performance. *Solar Energy* 83, 5 (2009), 614–624. doi:10.1016/j.solener.2008.10.008

[24] Swetha PT Srinivasan and Umesh Bellur. 2015. Watttime: novel system power model and completion time model for dvfs-enabled servers. In *2015 IEEE 21st International Conference on Parallel and Distributed Systems (ICPADS).* IEEE, 448–455.

[25] Thanawut Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2024. On the limitations of carbon-aware temporal and spatial workload shifting in the cloud. In *Proceedings of the 19th European Conference on Computer Systems (EuroSys).*

[26] Jan-Frederik Uhlenkamp, Jannicke Baalsrud Hauge, Eike Broda, Michael Lütjen, Michael Freitag, and Klaus-Dieter Thoben. 2022. Digital twins: A maturity model for their classification and evaluation. *IEEE Access* 10 (2022), 69605–69635.

[27] WattTime. 2025. Marginal Emissions API. https://www.watttime.org. Accessed: 2025.

[28] World Resources Institute. 2022. Global Power Plant Database. https://datasets.wri.org/datasets/global-power-plant-database. Accessed: 2025.

[29] Bobby Xiong, Davide Fioriti, Fabian Neumann, Iegor Riepin, and Tom Brown. 2024. Prebuilt Electricity Network for PyPSA-Eur based on OpenStreetMap Data. Zenodo dataset. doi:10.5281/zenodo.14144752 Version 0.6, accessed 2026-05-18.

[30] Bin Xiong, Davide Fioriti, Fabian Neumann, Iegor Riepin, and Tom Brown. 2025. Modelling the high-voltage grid using open data for Europe and beyond. *Scientific Data* 12, 277 (2025).

[31] Laura Yin, Logan Byers, Laura Malaguzzi Valeri, and Johannes Friedrich. 2020. *Estimating power plant generation in the Global Power Plant Database.* Technical Note. World Resources Institute, Washington, DC.

[32] Matteo Zanotto, Gabriele Padovani, Giovanni Iacca, and Sandro Fiore. 2026. Carbon Intensity Forecasting with On-Site Renewables for Sustainable AI Workload Scheduling. In *Proceedings of the 2nd International Workshop on Systems and Methods for Sustainable Large-Scale AI (GreenSys) (GreenSys '26).* Association for Computing Machinery, New York, NY, USA, 13–18. doi:10.1145/3802973.3804455