Türkiye Grid Foundation Model

A Reproducible Data Audit and Open Research Infrastructure for Multi-Task Electricity Forecasting

Faramarz Kowsari

Independent Researcher, Istanbul, Türkiye

ORCID: 0000-0003-1692-0453

Preprint - August 2026

Abstract

Electricity demand, renewable generation, and day-ahead market price are not independent signals: they arise within the same physical and market system and may therefore share useful temporal structure. That observation, however, does not imply that a shared multi-task model will necessarily improve forecasting. The longer-term question of the Türkiye Grid Foundation Model project is whether shared temporal representations can produce positive transfer across electricity consumption, renewable generation, and day-ahead Market Clearing Price (MCP; PTF in Turkish terminology), and under which regimes such sharing may instead produce negative transfer. This preprint addresses an earlier methodological requirement: whether the empirical data layer is sufficiently well defined, complete, aligned, and provenance-preserved to support a later model-comparison study.

Before the first empirical acquisition, the audit interval was committed as a scope lock from 2024-01-01T00:00:00+03:00 through 2024-01-31T23:00:00+03:00, corresponding to 744 expected hourly timestamps. Three official EPİAŞ Transparency Platform datasets were examined: real-time consumption, source-level real-time generation, and day-ahead MCP/PTF. Each dataset contained 744 raw rows and 744 unique hourly timestamps, with zero duplicates, zero missing expected hours, zero key numeric nulls, and zero observations outside the locked interval. Timestamp alignment yielded 744/744 complete three-target rows.

The initial audit also identified a generation-accounting residual when the reported total was compared with the listed generation-source sum. A corrective audit made no new EPİAŞ request and reused the exact locked January 2024 generation snapshot. When the `importExport` field was added to the listed source sum, the reported generation total reconciled to numerical precision across all 744 rows. Two solar observations of -0.01 MWh were documented and preserved without deletion or clipping. The contribution of this preprint is therefore deliberately bounded: a reproducible empirical data audit, an explicit evidence firewall between exploratory and future confirmatory work, and an open research infrastructure for subsequent multi-task forecasting and uncertainty-calibration experiments. No claim of forecasting superiority, state-of-the-art performance, or confirmatory model advantage is made. The confirmatory holdout remains undefined and untouched.

Keywords: electricity forecasting; EPİAŞ; data audit; reproducible research; multi-task learning; time-series forecasting; renewable generation; day-ahead electricity market; uncertainty quantification; Türkiye; open science.

1. Introduction

Electricity systems are composed of interacting physical and market processes rather than isolated time series. Demand, generation mix, renewable output, and day-ahead prices are observed on a common temporal axis and may contain overlapping information about system state. This motivates a natural modeling question: can information learned from one task improve another task when the targets are forecast jointly rather than independently?

The answer is not predetermined. Shared representations can create positive transfer when tasks expose stable common structure, but they can also create negative transfer when the relationships among tasks vary across market or system regimes. A model that improves average error in one period may fail during price shocks, rapid renewable ramps, or atypical demand conditions. For this reason, the Türkiye Grid Foundation Model treats multi-task transfer as an empirical hypothesis rather than a design assumption.

Short-term load forecasting has a long history in power-system planning and operations, and recent work has increasingly emphasized deep-learning models, uncertainty-aware forecasting, and practical feasibility [5]. The initial shared architecture in this project uses a Transformer encoder [3], while later uncertainty calibration may incorporate conformal methods [4]. These choices define a planned experimental framework; they are not themselves findings of the present paper.

Time-series evaluation also requires stronger safeguards than random train-test splitting. Daily and weekly periodicity, seasonality, market regime change, and operational shifts can allow random shuffling to mix information from different temporal regimes. The project therefore requires chronological evaluation, training-only fitting of learned preprocessing steps, and a strict separation between exploratory development and later confirmatory inference.

Before any model comparison is scientifically interpretable, however, the data layer must be auditable. Missing timestamps, duplicates, timezone errors, misunderstood source fields, or incorrect accounting conventions can invalidate conclusions even when a forecasting model reports a small error. In this project, the data audit is therefore part of the methodological design rather than a preliminary housekeeping step.

EPİAŞ provides official Transparency Platform web services for the Turkish electricity system, including real-time consumption, source-level generation, and day-ahead market-clearing prices [1]. The platform's credential-based access system was formally introduced for the new Transparency Platform login workflow in August 2024 [2]. The availability of these related hourly signals in a common official infrastructure makes them suitable for a reproducible multi-task research program, provided that their integrity and semantics are first checked explicitly.

This paper reports that first empirical checkpoint. Its contributions are fourfold: it documents a pre-committed audit interval; establishes complete hourly coverage and alignment for the three selected datasets within that interval; resolves an observed generation-accounting discrepancy by testing the role of `importExport` on the preserved snapshot; and connects the empirical record to a public open-science chain spanning GitHub, Zenodo, and OSF. The paper intentionally stops short of any claim about the predictive superiority of the shared model.

2. Research Design and Evidence Boundary

The planned forecasting problem uses a one-week hourly context window (168 hours) to predict three targets: electricity consumption, renewable generation, and day-ahead MCP/PTF. Candidate horizons are +1 h, +6 h, and +24 h. The initial shared model consists of a linear input projection, sinusoidal positional encoding, and a Transformer encoder feeding multi-horizon output heads. The term `Foundation Model` is the project name for a shared research platform and representation-learning program; it is not a claim that the current compact model has the scale or generality of large foundation models used in other domains.

The research lifecycle is divided into stages to prevent exploratory choices from being retrospectively described as confirmatory. Stage 1 is the bounded real-data audit reported here. Stage 2 expands the exploratory historical availability audit and benchmarks strong baselines while measuring runtime and compute requirements. Stage 3 freezes the final target, horizon, comparator set, preprocessing rules, seed policy, metrics, exclusions, and inferential procedure. Stage 4 creates a separate confirmatory OSF preregistration. Only after those steps will a confirmatory holdout be defined and opened for final evaluation.

Stage	Current status	Purpose
Bounded real-data audit	`Completed`	`Coverage, schema, alignment, accounting, provenance`
Exploratory modeling and compute audit	`Next`	`Baselines, feasibility, runtime and model selection`
Protocol freeze	`Not completed`	`Freeze target, horizon, models, seeds, metrics and inference`
Confirmatory OSF preregistration	`Not completed`	`Time-stamp the final confirmatory design`
Confirmatory holdout evaluation	`Not completed`	`Untouched final evaluation`

Table 1. Evidence-preserving research stages. The current preprint reports only the completed bounded data-audit stage.

At the time of this preprint, the confirmatory holdout remains `UNDEFINED_AND_NOT_REQUESTED`. This boundary is deliberate. It preserves the possibility that later results may show positive transfer, negative transfer, a tie, or an inconclusive difference without post-hoc tuning to the final evaluation period.

3. Data Sources and Audit Methods

3.1 Official EPİAŞ source contract

The audit uses official EPİAŞ Transparency Platform electricity web services [1]. Authentication is obtained through the EPİAŞ CAS/TGT service. The three endpoints used in the first empirical audit are:

```
CAS/TGT: https://giris.epias.com.tr/cas/v1/tickets
Real-time consumption: POST /v1/consumption/data/realtime-consumption
Real-time generation: POST /v1/generation/data/realtime-generation
Day-ahead MCP/PTF: POST /v1/markets/dam/data/mcp
```

EPİAŞ technical documentation identifies the real-time consumption service as hourly realized consumption, the real-time generation service as hourly generation by source, and the MCP/PTF service as the hourly day-ahead market-clearing price [1].

3.2 Pre-committed scope lock

Before the first empirical acquisition, the repository committed a dedicated scope-lock record. The audit window was fixed from 2024-01-01T00:00:00+03:00 to 2024-01-31T23:00:00+03:00, for 744 expected hourly timestamps. The first-audit runner did not expose command-line date overrides. The retrieval timestamp recorded by the audit was 2026-08-25T13:44:01.208005+00:00. No final confirmatory holdout dates were selected, and the runner requested no observations outside the locked January 2024 interval.

3.3 Third-party raw-data policy

Raw EPİAŞ records are third-party source data and are not redistributed by the repository. Raw and processed records remain local and are excluded from Git by default. The public research record preserves source endpoints, retrieval metadata, schema observations, aggregate quality summaries, and cryptographic hashes of local snapshots. Authentication credentials are not stored in the repository.

3.4 Temporal coverage and cross-dataset alignment

For each dataset, the audit compared raw row count, unique hourly timestamps, and the expected hourly grid. Duplicate timestamps, missing expected hours, and observations outside the locked interval were counted explicitly. Temporal coverage was defined as:

```
Coverage (%) = (N_unique / N_expected) x 100
```

The three datasets were then aligned by hourly timestamp. A timestamp was considered a complete three-target row only when consumption, generation, and MCP/PTF were all present. The alignment rate was defined as:

```
Alignment rate (%) = (N_complete_three_target_rows / N_expected) x 100
```

3.5 Schema, timezone, and numeric integrity

The audit recorded observed field names, timezone suffixes, and the maximum gap between consecutive timestamps. The generation schema included source fields such as `fueloil`, `blackCoal`, `lignite`, `geothermal`, `naturalGas`, `river`, `dammedHydro`, `lng`, `biomass`, `naphta`, `importCoal`, `asphaltiteCoal`, `wind`, `sun`, and `wasteheat`, together with `importExport` and reported `total`. All three datasets used the `+03:00` timezone suffix, and the maximum timestamp gap observed in each was 1.0 hour.

For key numeric variables, the audit calculated count, mean, standard deviation, minimum, 5th percentile, median, 95th percentile, and maximum. It also counted key nulls and basic anomaly indicators. Anomaly flagging was intentionally separated from deletion: observations were not removed merely because they were unusual.

3.6 Generation-accounting diagnostic

The first audit compared reported total generation with the sum of listed source-level generation fields. Because the `importExport` field was present in the schema, the corrective audit tested three accounting conventions on the exact preserved generation snapshot rather than assuming the sign of `importExport` in advance.

```
r_t^(0) = G_reported,t - Sum_j G_j,t
r_t^(+) = G_reported,t - (Sum_j G_j,t + IE_t)
r_t^(-) = G_reported,t - (Sum_j G_j,t - IE_t)
```

For each convention, the audit calculated the mean residual, mean absolute residual, 95th percentile of the absolute residual, and maximum absolute residual. Mean absolute residual was defined as:

```
MAE_r = (1 / N) x Sum_t |r_t|
```

3.7 Cryptographic provenance

SHA-256 hashes were recorded for each local raw snapshot and for the processed aligned CSV. These hashes do not replace the source data and do not independently reconstruct them; they provide an identity check for determining whether a later analysis uses the same preserved snapshot.

4. Results

4.1 Hourly coverage and alignment

Dataset	Raw rows	Unique hours	Duplicates	Missing	Coverage	Key nulls	Outside scope
Consumption	744	744	0	0	100.000%	0	0
Generation	744	744	0	0	100.000%	0	0
MCP/PTF	744	744	0	0	100.000%	0	0

Table 2. Hourly coverage and integrity within the locked January 2024 audit interval.

All three datasets covered the complete expected hourly grid. Each contained 744 raw rows and 744 unique hourly timestamps, with zero duplicates, zero missing expected hours, zero key numeric nulls, and zero observations outside scope. Alignment by timestamp produced 744/744 complete three-target rows, corresponding to a 100.000% complete-alignment rate.

4.2 Descriptive numeric audit

Variable	N	Mean	SD	Min	P05	Median	P95	Max
Consumption (MWh)	744	38,873.309	5,767.913	24,610.500	30,137.229	39,592.630	46,900.831	48,394.390
Renewable generation (MWh)	744	19,081.383	3,440.859	10,398.190	12,745.366	19,116.315	24,406.332	26,093.290
Total generation (MWh)	744	37,898.719	5,400.053	24,536.760	29,426.151	38,574.885	45,589.784	47,992.460
Renewable share	744	0.505	0.068	0.350	0.396	0.503	0.625	0.673
MCP/PTF (TL/MWh)	744	1,942.905	626.953	102.260	900.000	2,180.980	2,689.000	2,700.000

Table 3. Descriptive statistics for the primary audited variables. These are data-quality descriptions, not forecasting metrics.

The primary-variable IQR diagnostic flagged zero observations for consumption, the renewable aggregate, and MCP/PTF. The descriptive Pearson correlations were 0.6755 between consumption and renewable generation, 0.6778 between consumption and MCP/PTF, and 0.1896 between renewable generation and MCP/PTF. These correlations are exploratory, non-causal, and not confirmatory evidence of cross-task transfer.

4.3 Corrective generation-accounting audit

Convention	N	Mean residual	MAE residual	P95 abs.	Max abs.
Exclude importExport	744	-88.345793	121.439449	274.696500	380.370000
Add importExport	744	-0.000000	0.000000	0.000000	0.000000
Subtract importExport	744	-176.691586	242.878898	549.393000	760.740000

Table 4. Corrective comparison of three `importExport` accounting conventions on the same locked generation snapshot.

The corrective audit made zero new EPİAŞ requests. Excluding `importExport` produced a mean absolute residual of 121.439449 MWh, while subtracting it increased the mean absolute residual to 242.878898 MWh. Adding `importExport` produced zero residual to the displayed numerical precision across all 744 rows. The supported interpretation is therefore limited to this locked snapshot: the observed January 2024 accounting convention was consistent with reported total generation equaling the listed source sum plus `importExport`, to numerical precision. This finding is not generalized to all historical EPİAŞ periods without additional audit.

4.4 Negative solar observations

The generation snapshot contained two negative `sun` observations, both equal to -0.01 MWh, at 2024-01-04T18:00:00+03:00 and 2024-01-13T18:00:00+03:00. Their total absolute magnitude was 0.02 MWh. Neither value was deleted, clipped, or replaced. The purpose of the audit at this stage was to document source behavior rather than silently normalize unusual values.

4.5 Snapshot identity

Dataset	Local snapshot	SHA-256
Consumption	`consumption.json`	49a528ee6ffbd3c1a2717cddc7bbbb311a5c4a52a6a0379ea05dc1f2fc9ddc18
Generation	`generation.json`	63afec4aed0c73bef2b374ae356e3e1914d85be23f167a90b7db8682fb7ff344
MCP/PTF	`mcp.json`	1853a77aa7ef96d7ecbc7bfef0896e151954251e2dae54318b06043fde452513

Table 5. SHA-256 identifiers for the three local raw snapshots.

`Processed aligned CSV SHA-256: d0cc98d066fafcd71868b876a2b44dc292a5c597f6c0df3661a0e4620741b512`

5. Discussion

The first data gate supports continued exploratory development, but its scope should not be overstated. Within the locked January 2024 interval, the three selected EPİAŞ datasets were complete on the expected hourly grid, aligned cleanly by timestamp, and free of key numeric nulls. This establishes a defensible empirical starting point for the next stage of the project.

The generation-accounting result illustrates why data semantics must be audited before modeling. The initial residual could have been misclassified as a generic data-quality defect. Re-examining the preserved snapshot under alternative `importExport` conventions instead showed that the discrepancy was consistent with a specific accounting interpretation. This distinction matters because arbitrary correction of the raw source fields would have weakened reproducibility.

The two -0.01 MWh solar observations are numerically small, but their preservation reflects an important methodological rule. Data cleaning choices can become hidden degrees of freedom if they are made silently after inspection. Recording the values without clipping keeps later preprocessing decisions explicit and auditable.

Most importantly, the current findings say nothing about whether the shared Transformer will outperform strong baselines. The observed correlations are not evidence of transfer, and the one-month audit is not a forecasting benchmark. The paper's contribution is the integrity of the empirical starting point and the explicit boundary around what has and has not yet been tested.

6. Reproducibility and Open-Science Record

The project maintains complementary public records rather than treating a single platform as the entire evidence chain. GitHub is the living development repository. Zenodo archives citable software snapshots, and OSF records the research lifecycle and evidence boundary. The v0.2.1 software snapshot is archived on Zenodo under DOI 10.5281/zenodo.22102736, while the Zenodo Concept DOI is 10.5281/zenodo.22102735 [6]. The accepted public OSF Open-Ended Registration is identified by DOI 10.17605/OSF.IO/FMCYQ [7].

Record	Role	Persistent identifier / location
GitHub	Living code, reports, protocol and documentation	github.com/FaramarzKowsari/turkiye-grid-foundation-model
Project website	Public research interface	faramarzkowsari.github.io/turkiye-grid-foundation-model/
Zenodo v0.2.1	Exact archived software snapshot	10.5281/zenodo.22102736
Zenodo Concept DOI	Persistent record across software versions	10.5281/zenodo.22102735
OSF Open-Ended Registration	Lifecycle and evidence-boundary record	10.17605/OSF.IO/FMCYQ
Associated OSF project	Related OSF project space	osf.io/tz5pw

Table 6. Public open-science record chain for the project.

The OSF Open-Ended Registration is not the future confirmatory preregistration. It time-stamps the current research state after the initial empirical data audit and documents the intended separation between exploratory development and later confirmatory evaluation. A separate confirmatory registration will be created only after the experimental protocol is frozen.

7. Limitations and Next Confirmatory Stage

The present audit covers one calendar month. It therefore does not establish year-round completeness, long-run schema stability, or representativeness across all demand, renewable, and price regimes. January 2024 cannot be assumed to represent summer demand, extreme-heat periods, large market shocks, unusual renewable conditions, or structural market changes.

No forecasting benchmark is reported in this paper. The final model set, primary target and horizon, seed policy, preprocessing specification, exclusion rules, and statistical inference procedure are not yet frozen. The observed correlation structure must not be interpreted as proof that joint learning will help. Similarly, the generation-accounting convention identified here is limited to the audited snapshot and must be rechecked when the historical scope is expanded.

The next empirical stage will extend the exploratory historical-availability audit on the training and validation side, benchmark strong baselines such as persistence and seasonal-naive forecasting, and measure training time, inference time, storage requirements, and total compute per model and seed. These measurements are required before freezing a feasible confirmatory protocol. After that stage, the project will fix the primary estimand, comparator set, seeds, metrics, exclusions, preprocessing, inferential procedure, and holdout definition in a separate confirmatory OSF preregistration. Only then will the confirmatory holdout be opened.

8. Conclusion

This preprint reports the first empirical result of the Türkiye Grid Foundation Model at the level that has actually been tested: the integrity and reproducibility of the initial data layer. In the locked January 2024 interval, consumption, source-level generation, and MCP/PTF each covered all 744 expected hourly timestamps, with no duplicates, no missing expected hours, no key numeric nulls, and complete timestamp alignment across the three targets.

The initial generation-total discrepancy was resolved through a corrective accounting audit rather than by altering the source data. On the same preserved snapshot, adding `importExport` to the listed generation-source sum reconciled reported total generation to numerical precision. Two small negative solar observations were documented and retained without clipping.

The project therefore has an auditable empirical foundation for the next exploratory modeling stage. This paper is not evidence that a Transformer is more accurate than isolated-task models, nor that multi-task learning improves Turkish electricity forecasting. Those remain open empirical questions. The purpose of the current record is to show how the data, provenance, accounting semantics, and evidence boundary were established before those questions are tested.

Data and Code Availability

Source code, audit reports, protocol documents, and public research records are available from the project repository and related persistent records. Raw third-party EPİAŞ records are not redistributed by the repository; reproducibility is supported through official retrieval instructions, source endpoints, audit metadata, and cryptographic snapshot identifiers.

GitHub repository: https://github.com/FaramarzKowsari/turkiye-grid-foundation-model

Project website: https://faramarzkowsari.github.io/turkiye-grid-foundation-model/

Zenodo v0.2.1 DOI: https://doi.org/10.5281/zenodo.22102736

Zenodo Concept DOI: https://doi.org/10.5281/zenodo.22102735

OSF Registration DOI: https://doi.org/10.17605/OSF.IO/FMCYQ

Associated OSF project: https://osf.io/tz5pw

Competing Interests

The author declares no known financial, commercial, or organizational competing interests that could reasonably be perceived as influencing the work reported in this preprint.

Author Contributions

Faramarz Kowsari: conceptualization, methodology, software, investigation, data curation, reproducibility design, writing - original manuscript preparation, and writing - review and editing.

AI-Assisted Language and Editorial Disclosure

OpenAI ChatGPT was used for English-language translation of an author-prepared Persian manuscript and for copy-editing, formatting, and organization of existing project documentation. The empirical values reported in this paper originate from the project's preserved audit outputs; no synthetic empirical results were introduced. No AI system is listed as an author. The human author remains responsible for source verification, numerical accuracy, interpretation, and approval of the final wording.

References

[1] EPİAŞ. Transparency Electricity Services Technical Documentation (Şeffaflık Elektrik Servisleri). Energy Exchange Istanbul. Available at: https://seffaflik.epias.com.tr/electricity-service/technical/tr/index.html (accessed 27 August 2026).

[2] EPİAŞ. The Registration for the Transparency Platform is Now Open. Electricity Market Transparency Platform. 14 August 2024. Available at: https://www.epias.com.tr/en/announcements/market/electricity/the-registration-for-the-transparency-platform-is-now-open/

[3] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I. Attention Is All You Need. Advances in Neural Information Processing Systems. 2017;30:5998-6008. Available at: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[4] Angelopoulos AN, Bates S. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. arXiv preprint arXiv:2107.07511. 2021. Available at: https://arxiv.org/abs/2107.07511

[5] Eren Y, Kucukdemiral IB. A comprehensive review on deep learning approaches for short-term load forecasting. Renewable and Sustainable Energy Reviews. 2024;189:114031, 1-60. doi:10.1016/j.rser.2023.114031.

[6] Kowsari F. Türkiye Grid Foundation Model v0.2.1. Zenodo; 2026. doi:10.5281/zenodo.22102736. Concept DOI: 10.5281/zenodo.22102735.

[7] Kowsari F. Türkiye Grid Foundation Model: Cross-Task Transfer And Calibrated Uncertainty For Multi-Horizon Forecasting In Türkiye's Electricity System. OSF Registries; 2026. doi:10.17605/OSF.IO/FMCYQ.