ISBN: 978-628-96613-3-0. ISSN: 2414-6390. Digital Object Identifier: https://dx.doi.org/10.18687/LACCEI2026.1.1.1391

Data-Driven Hemodialysis Demand Estimation: Comparing Patient Grouping Alternatives for Weekly EPO Requirements and Dialysis Hours

Fabricio Smith[1]; Araceli Nuñez[2]; Andrea Cristaldo[3]; Eduardo Redondo[4]
[1,2,3,4]Facultad de Ingeniería, Universidad Nacional de Asunción, Paraguay
evaliente@fiuna.edu.py, amnunez@fiuna.edu.py, andrea.cristaldo@fiuna.edu.py, eredondo@ing.una.py

Abstract*– This study develops a data-driven framework to anticipate hemodialysis (HD) workload and erythropoietin (EPO) requirements in Paraguay ' s public healthcare system using routinely collected service records from a national nephrology center (2021 – 2024; n = 2,592 patients). The objective is to support operational planning of critical resources by combining patient grouping strategies with simple univariate forecasting approaches. After systematic data cleaning and harmonization, temporal consistency adjustments were applied using institutional year-to-year statistics to preserve the continuity of the demand series. We compared empirical stratification with unsupervised grouping methods (PAM, hierarchical clustering, and DBSCAN), and evaluated several forecasting rules (naïve, Holt, linear regression, and weighted moving average). Given the short annual history, a single-year holdout (2024) was used to consistently screen configurations. Under this evaluation protocol, empirical stratification coupled with a weighted moving average achieved the lowest weighted mean absolute errors: 10.79 patients, 126.47 HD hours/week, and 116,938 IU/week of EPO. The selected configuration was then used to construct uncertainty bands for 2025, projecting 6,444 – 8,115 HD hours/week and 4.43 – 5.95 million IU/week of EPO to inform resource planning.

Keywords*– Demand anticipation; Patient grouping; Unsupervised clustering; Erythropoietin (EPO); Hemodialysis; Healthcare resource planning.

I. Introduction

Chronic kidney disease (CKD) is a growing public health burden worldwide, affecting roughly one in ten adults and exhibiting sustained increases over recent decades, with major implications for health-system capacity and spending [1], [2]. Global burden assessments highlight both the magnitude of CKD and its upward trajectory across regions, with downstream growth in the need for kidney replacement therapies (KRT) as patients progress to advanced stages [1]. From a clinical standpoint, CKD is typically defined and staged according to kidney function decline (e.g., estimated glomerular filtration rate), and its progression is shaped by metabolic and cardiovascular risk profiles as well as access-to-care constraints [3]. International policy discussions have also positioned kidney disease within broader health and development agendas, emphasizing prevention and equitable access to effective treatment modalities [4].

When conservative management is insufficient, advanced CKD (including kidney failure) requires KRT—hemodialysis (HD), peritoneal dialysis, or transplantation—each with distinct infrastructure, workforce, and supply-chain requirements [5], [6]. In practice, HD remains highly prevalent in many settings and imposes recurrent operational demands: fixed treatment schedules, machine time, nursing time, vascular access management, dialysate preparation, and continuous procurement of consumables and pharmaceuticals [5], [7]. Adequacy targets and prescription elements (e.g., blood flow, ultrafiltration, dialysate composition, session time and frequency) are central to routine operations and are commonly framed using measures such as Kt/V [7], [8]. Moreover, alternative regimens—such as increased treatment frequency—can alter utilization patterns and intensify resource requirements, which reinforces the need for forward-looking planning tools that can translate service activity into capacity and procurement decisions [9].

In Latin America, HD is the dominant KRT modality; regional syntheses report that a substantial majority of prevalent KRT patients are treated with HD, although with marked between-country heterogeneity and persistent constraints in infrastructure and access [10]. These characteristics make operational planning especially sensitive to demand fluctuations, patient mix, and supply reliability. Beyond the dialysis session itself, HD safety depends on strict technical standards—particularly water quality—because large volumes of treated water are required to prepare dialysate and any deviation can have clinical consequences and operational disruptions [11]. Empirical assessments from dialysis units in the region further document that physical-chemical and microbiological parameters of water and dialysate require systematic monitoring, with potential noncompliance risks that can translate into unplanned downtime and additional costs [12].

A second major operational driver in CKD care is anemia management. Reduced endogenous erythropoietin production is a key mechanism in CKD-related anemia, and erythropoiesis-stimulating agents such as epoetin alfa are widely used in clinical practice [13]. Seminal economic evidence shows that recombinant human erythropoietin can change transfusion needs and cost structures in dialysis populations, underscoring its relevance as both a clinical and budgetary item [14]. More recent analyses continue to associate anemia-related patterns—such as hemoglobin stability—to healthcare resource utilization and costs among dialysis-dependent patients [15]. Together, these results justify treating erythropoietin (EPO) consumption not merely as a clinical variable, but as a planning-relevant demand stream that should be anticipated alongside HD hours.

Within Paraguay, public provision of HD faces additional structural challenges related to geographic concentration and the logistics of expanding coverage. National institutional reporting has documented a persistent concentration of dialysis centers in the Asunción/Central area, alongside a broader evolution of center distribution over time [16]. Complementary national reporting through dialysis and transplant registries provides epidemiologic and service-coverage context that is essential for interpreting demand signals and their policy implications, including multi-year series describing KRT coverage and patient counts [17]–[19]. In such settings, planning decisions—budgeting, procurement, staffing, and short-term capacity sizing—benefit from quantitative data-driven methods that can leverage routinely collected service records to produce transparent, reproducible projections.

Despite the strong operational motivation, the integration of patient grouping and short-horizon demand forecasting for HD workload and EPO consumption remains underdeveloped in many real-world public systems, particularly where clinical covariates are incomplete or inconsistently recorded. Practical overviews of dialysis care emphasize the complexity of modality choice, indications, and operational requirements, but do not by themselves provide an implementable forecasting workflow for resource planning at system level [20]. This gap motivates a data-driven approach that prioritizes (i) segmentation strategies that remain stable and interpretable under real-world data constraints, and (ii) forecasting methods that can be validated out-of-sample for planning horizons relevant to procurement and capacity decisions.

Accordingly, this study develops and evaluates a structured workflow to anticipate near-term demand for HD hours/week and EPO consumption using routinely collected service records from Paraguay. Two research questions are addressed: (i) which patient grouping strategies—empirical stratification versus unsupervised clustering—provide stable and useful segmentation for projecting HD hours/week and EPO consumption, and (ii) which univariate forecasting models yield lower out-of-sample error for the 2025 planning horizon. Using harmonized yearly demand series, we compare alternative grouping–forecasting configurations and produce planning-oriented outputs: projected totals and uncertainty bands for HD hours/week and EPO units. By grounding the analysis in national institutional reporting and established clinical/operational standards, the contribution is an evidence-based, reproducible pathway from service records to actionable resource-planning estimates in a public health setting [5], [7], [8], [16]–[19].

The remainder of the paper is organized as follows. Section II describes the data source, preprocessing and temporal harmonization steps, the alternative patient grouping strategies, the univariate forecasting models, and the out-of-sample evaluation protocol, including how projections are generated for planning use. Section III reports the descriptive characterization of the cohort, compares grouping outputs, presents forecasting performance results under the holdout design, and summarizes the selected configuration and its demand estimates. Section IV discusses the main methodological and operational implications, highlights limitations, and outlines directions for extension and validation in similar settings.

II. Methodology

This section describes the data foundation, the preprocessing steps used to build consistent demand series, the alternative patient-grouping strategies, the forecasting models and their parameterization, the out-of-sample evaluation protocol, and the procedure used to generate 2025 demand estimates and uncertainty ranges for planning purposes.

A. Data source, study window, and target variables

The study uses routinely collected clinical-administrative service records from a national public nephrology provider in Paraguay over the 2021–2024 period. The dataset contains patient identifiers and a small set of stable descriptors (e.g., age and sex), together with operational variables directly linked to resource planning, notably hemodialysis (HD) workload and erythropoietin (EPO) requirements.

Three demand outcomes were defined as annual snapshots of weekly planning indicators: (i) Yearly patient volume P_t: number of unique patients recorded in year t. (ii) HD workload intensity H_t: total HD hours/week required by the cohort recorded in year t, obtained by summing each patient's prescribed weekly HD hours within that year. (iii) EPO requirement intensity E_t: total EPO IU/week required by the cohort recorded in year t, obtained by summing each patient's prescribed weekly EPO dose within that year. These indicators are designed for operational planning (capacity sizing and drug procurement) and should not be interpreted as annual totals.

B. Data curation and construction of consistent demand series

Operational datasets typically include duplicated entries, inconsistent coding practices, and time-varying recording intensity. Because forecasting models are sensitive to artificial discontinuities, a structured curation workflow was applied to preserve the underlying temporal signal while ensuring full reproducibility.

1. **Harmonization of identifiers and categorical variables.** Patient identifiers were standardized to ensure stable linkage over time. Categorical fields (sex; dialysis modality/type when used) were mapped to a consistent set of values to avoid spurious group splits.

2. **Uniqueness rules and patient counting.** Patient volume P_t was computed using a unique-patient-per-year rule to prevent multiple records for the same individual inflating counts (and the same convention was applied at the segment level).

3. **Aggregation to annual snapshots of weekly indicators.** For each year t, cohort-level H_t (hours/week) and E_t (IU/week) were constructed by summing patient-level weekly prescriptions / requirements among patients recorded in that year. Importantly, the analysis does not multiply by 52 to obtain annual totals; instead, it forecasts the weekly requirement indicator associated with each year's cohort.

4. **Temporal reconciliation for continuity.** When an inter-annual coverage discontinuity was identified, a deterministic correction was applied to the affected year to restore continuity before model comparison. In practice, this consisted of correcting the patient count and coherently rescaling the cohort-level weekly indicators using the observed intensity structure for that year, and propagating the correction consistently to grouped series so that adjusted segment totals sum to the adjusted aggregate. The rescaling preserves within-year relative intensities while correcting for documented coverage gaps. This adjustment primarily affected the year 2023, where the observed patient count was lower than adjacent years; the correction resulted in a proportional rescaling of the demand indicators to restore continuity.

5. **Feature preparation for grouping methods.** For grouping strategies requiring a patient feature matrix (unsupervised clustering), we constructed a patient-level representation that combines stable descriptors and resource-use intensities derived from the curated period. Continuous features were standardized to avoid scale dominance in distance computations, and categorical features were encoded in a way compatible with mixed-type dissimilarities when applicable.

C. Patient grouping strategies and rationale

Forecasting in heterogeneous patient populations can be improved by modeling segments with more homogeneous resource-use patterns. However, in operational settings, grouping methods must satisfy three constraints: interpretability (to support action), stability (to support consistent procurement rules), and feasibility (to avoid segments too small to model). We therefore compared:

1. **Empirical stratification (interpretable baseline).** Patients were assigned to groups based on a cross-classification of age category and sex. This choice is justified by: (i) strong operational interpretability, (ii) low sensitivity to noise in poorly recorded clinical markers, and (iii) stable group membership over time (except for aging across categories, which is predictable). The resulting strata are straightforward to communicate and maintain as part of institutional planning.

2. **Unsupervised clustering (data-driven alternative).** Three commonly used unsupervised methods representing distinct clustering principles were evaluated:

   - **PAM (Partitioning Around Medoids):** robust to outliers and suitable for dissimilarity matrices; it produces representative medoids and tends to be stable under moderate noise.

   - **Hierarchical clustering (agglomerative, Ward-type criterion):** builds a nested structure that can be cut at different granularities, enabling a controlled trade-off between segment count and within-segment cohesion.

   - **DBSCAN:** identifies dense regions and can isolate atypical patients as noise, which is relevant in clinical-administrative data where rare profiles may exist.

The purpose of including these methods is not to assume that clustering is inherently superior, but to test whether data-driven groupings yield demand series that are more forecastable than policy-relevant empirical strata under real-world recording constraints.

Parameter selection for clustering

To prevent arbitrariness in grouping outcomes, clustering parameters were selected using objective, reproducible rules:

- For **PAM** and **hierarchical clustering**, the number of clusters k was chosen from a candidate set by maximizing an internal validity criterion (e.g., average silhouette) subject to feasibility constraints such as a minimum cluster size and the avoidance of degenerate solutions.

- For **DBSCAN**, ε and MinPts were selected through a systematic grid search guided by density diagnostics and internal validity criteria, while enforcing the same feasibility constraints. Patients labeled as noise were explicitly tracked and accounted for in downstream aggregation adjustments.

These choices ensure that clustering comparisons reflect principled parameterization rather than post-hoc tuning to match desired outcomes.

D. Forecasting models and parameterization

Forecasting was conducted independently for each segment g and each outcome $\{P, H, E\}$, using univariate models that are appropriate when the primary objective is reliable short horizon planning with limited explanatory covariates. This also ensures that any observed performance differences can be attributed to grouping strategy and time-series structure, rather than complex multivariate modeling choices.

Four forecasting model families were evaluated:

1. **Naïve:** $\hat{Y}_{t+1} = Y_t$. This provides a strong baseline in operational series where persistence is common.

2. **Holt's linear method:** captures local trend without requiring strong seasonality assumptions; smoothing parameters are estimated automatically by minimizing in-sample loss on the training set.

3. **Linear time regression:** models systematic drift using a simple parametric form; coefficients are estimated by ordinary least squares on the training window.

4. **Weighted moving average (exponential form).** A parsimonious recency-weighted rule was implemented in exponential smoothing form, $\hat{Y}_{t+1} = \alpha Y_t + (1-\alpha)\hat{Y}_t$, with $\alpha \in (0,1)$. The smoothing parameter was selected through a small grid search to minimize absolute error under the single-year benchmarking design.

Parameter selection for forecasting

To avoid methodological gaps, parameters were defined through explicit, non-adaptive rules:

- For **Holt**, smoothing parameters were estimated via standard optimization on the training data for each segment-series, using a fixed objective and bounded parameter space.

- **For WMA (exponential form),** the smoothing parameter α was selected from a small, predefined grid by minimizing holdout absolute error under the fixed benchmarking protocol. This preserves comparability across segments and prevents ad hoc per-series tuning.

- For **regression**, the model form (with or without seasonal indicators) was fixed before evaluation; model selection was driven by test performance under the same protocol for all segments.

E. Out-of-sample evaluation design, weighting, and selection rule

Given the very short annual history, 2024 was used as a single-year benchmark to consistently screen configurations rather than as a definitive statistical validation. For each configuration (grouping strategy G + forecasting model M), models were fitted on the training window (2021–2023) and used to generate a one-step-ahead forecast for the held-out year 2024 for every segment g and outcome $Y \in \{P, H, E\}$.

Because the holdout window contains a single observation per segment-series at yearly resolution, forecast accuracy in 2024 is reported using (i) the absolute error in natural units and (ii) the symmetric mean absolute percentage error (sMAPE) to enable scale-aware comparison across outcomes. Segment-level errors were aggregated into global planning scores using nonnegative weights w_g (with $\sum_g w_g = 1$). In the implemented evaluation, weights were set proportional to the observed segment magnitude in the holdout year, i.e., $w_g = Y_{g,2024} / \sum_h Y_{h,2024}$, to reflect the relative contribution of each segment when summarizing errors.

$$\mathrm{WMAE}(Y) = \sum_g w_g \, | Y_{g,2024} - \hat{Y}_{g,2024} |$$

$$\mathrm{WsMAPE}(Y) = \sum_g w_g \, \frac{2 \, | Y_{g,2024} - \hat{Y}_{g,2024} |}{| Y_{g,2024} | + | \hat{Y}_{g,2024} | + \varepsilon}$$

where ε is a small constant to prevent division by zero. For each outcome, the preferred configuration was selected by minimizing $\mathrm{WMAE}(Y)$, using $\mathrm{WsMAPE}(Y)$ as a secondary criterion for tie-breaking and consistency checks.

F. 2025 projections and uncertainty ranges

After selecting the preferred configuration under the 2024 test protocol, the corresponding segment-level models were refitted using the full 2021–2024 history and generated forecasts for the 2025 planning year label (weekly indicators). Aggregate demand projections were obtained by summing segment forecasts:

$$\hat{Y}_t = \sum_g \hat{Y}_{g,t}, \, Y \in \{P, H, E\}$$

Handling exclusions and ensuring comparability

When DBSCAN labels patients as noise (excluded from any cluster), excluding their demand would artificially reduce projected totals. To ensure comparability across grouping strategies, totals were homogenized by re-incorporating patients not represented in the segment-level forecasting stage (including DBSCAN noise points and other rule-based exclusions) and assigning them the average weekly EPO dose (IU/week) and weekly HD hours (hours/week) observed in the corresponding analyzable set.

The resulting add-on demand was computed as the product between these averages and the number of missing patients and then added to the 2025 projected totals so that all scenarios represent the full patient population under every grouping approach.

Algorithm overview

Algorithm 1 summarizes the end-to-end workflow as implemented: data curation and temporal reconciliation; definition of patient groupings under empirical and unsupervised alternatives with principled parameter selection;

construction of segment-level annual snapshots of weekly demand indicators; fitting and evaluation of univariate forecasting rules using a single-year holdout benchmark (2024); selection of the preferred configuration; refitting on the full window; and generation of 2025 projections with coherent aggregation and planning-oriented uncertainty ranges.

Algorithm 1. Segmented demand anticipation workflow (HD hours/week and EPO IU/week).

Inputs: records R (2021–2024); outcomes $\{P, H, E\}$; grouping set $\mathcal{G}$; forecasting rules $\mathcal{F}$;
$\mathcal{T}_{train} = 2021–2023$; $\mathcal{T}_{hold} = 2024$; horizon $H = 1$ (2025);
parameter grids Π_G, Π_M; feasibility constraints Φ; weights w_g; ε.
Outputs: best configuration $C^*(Y)$; forecasts $\hat{Y}_{Y,t}$ (2025); bands $[L_{Y,t}, U_{Y,t}]$.
Curate R; build annual snapshots $\{P_t, H_t, E_t\}$, where H_t (hours/week) and E_t (IU/week) are cohort-level weekly indicators computed within each year t.
If a discontinuity is detected: reconcile totals and propagate coherently to segment series.
$\Omega \leftarrow \emptyset$
FOR $G \in \mathcal{G}$ **DO**
| Assign groups $g(p)$: strata rules if empirical; else Cluster(R, θ_G) with $\theta_G \in \Pi_G$ (validity + Φ)
| Build segment series $Y_{g,t}$ for $Y \in \{P, H, E\}$
| **FOR** $M \in \mathcal{F}$ **DO**
| | Select $\theta_M \in \Pi_M$ (model-specific, fixed rule)
| | **Benchmark on** $\mathcal{T}_{hold}$:
| | **FOR** $Y \in \{P, H, E\}$ **DO**
| | | Fit M on $\mathcal{T}_{train}$ for each g; forecast 2024; compute Errors(Y) using w_g and ε
| | **END FOR**
| | Refit on 2021–2024 and forecast 2025:
| | **FOR** $Y \in \{P, H, E\}$ **DO**
| | | Fit M for each g; forecast $\hat{Y}_{g,2025}$; aggregate $\hat{Y}_{Y,2025} \leftarrow \sum_g \hat{Y}_{g,2025}$
| | | If G yields noise/exclusions: re-incorporate excluded demand coherently (add-back rule)
| | **END FOR**
| | If (G, M) satisfies Φ: $\Omega \leftarrow \Omega \cup \{(G, M, \text{Errors}(\cdot), \hat{Y}_{(\cdot),2025})\}$
| **END FOR**
END FOR
FOR $Y \in \{P, H, E\}$ **DO**
| $C^*(Y) \leftarrow \arg\min_{(G,M)\in\Omega} \text{Errors}(Y)$ (primary: WMAE; secondary: WsMAPE)
| Set $\hat{Y}_{Y,2025} \leftarrow \hat{Y}^{C^*(Y)}_{Y,2025}$
| $L_{Y,2025} \leftarrow \min_{(G,M)\in\Omega} \hat{Y}^{(G,M)}_{Y,2025}$; $U_{Y,2025} \leftarrow \max_{(G,M)\in\Omega} \hat{Y}^{(G,M)}_{Y,2025}$
END FOR
RETURN $\{C^*(Y), \hat{Y}_{Y,2025}, [L_{Y,2025}, U_{Y,2025}]\}$

III. RESULTS

A. Cohort profile and operational heterogeneity (HD modality, HD hours/week, and weekly EPO requirements)

All workload and drug outcomes were reported as annual snapshots of cohort-level weekly requirements (hours/week; IU/week), not annual totals. Figure 1 depicts the annual evolution of the treated cohort size, providing context for the demand indicators analyzed in this study.

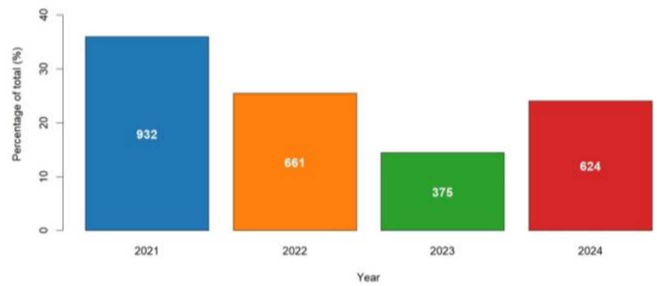

Fig 1 Annual evolution of the treated hemodialysis cohort in the INN registry (2021–2024): number of unique patients per year.

A key operational feature of the registry is that the cohort is overwhelmingly treated with in-center hemodialysis, with peritoneal dialysis representing a small but non-negligible fraction each year (Table 1).

TABLE I
DIALYSIS MODALITY BY YEAR (COUNTS AND %)

Year	Modality	N	%
2021	Hemodialysis	536	95.89%
2021	Peritoneal dialysis	23	4.11%
2022	Hemodialysis	623	94.25%
2022	Peritoneal dialysis	38	5.75%
2023	Hemodialysis	351	94.10%
2023	Peritoneal dialysis	22	5.90%
2024	Hemodialysis	586	94.98%
2024	Peritoneal dialysis	31	5.02%

Beyond modality, the dataset documents substantial dispersion in weekly HD time requirements. Table 2 summarizes the overall distribution of weekly HD hours and its year-by-year breakdown, showing stable central tendency (median = 12) with meaningful variation in dispersion across years. (Table 2 summarizes patients with non-missing weekly HD hours (HD modality), hence N < base cohort).

TABLE II
DISTRIBUTION OF WEEKLY HD HOURS (OVERALL AND BY YEAR)

Scope	N	Mean	Median	SD	Min	Max
Overall	1,897	11.07	12	2.88	4	30
2021	508	10.48	12	3.85	4	30
2022	538	11.27	12	2.63	4	28
2023	332	11.69	12	1.55	4	12
2024	519	11.06	12	2.55	4	12

To support an interpretable segmentation strategy under routine data constraints, the empirical grouping is defined as age group × sex (10 groups), with group sizes varying across years (Table 3).

TABLE III
EMPIRICAL STRATA (AGE × SEX): GROUP SIZES BY YEAR

Empirical group	2021	2022	2023	2024
18–29, Female	19	18	9	15
18–29, Male	21	19	10	12
30–44, Female	40	34	22	31
30–44, Male	53	41	30	40
45–59, Female	53	71	30	56
45–59, Male	80	102	45	97
60–74, Female	65	85	53	61

60–74, Male	118	112	73	97
≥75, Female	11	10	8	16
≥75, Male	18	22	19	19

Table 4 quantifies how this segmentation translates into different resource intensities. For paper concision, group averages for weekly HD hours and weekly EPO requirements were reported for the most recent year (2024).

TABLE IV
EMPIRICAL STRATA RESOURCE INTENSITY (2024): MEAN WEEKLY HD HOURS AND MEAN WEEKLY EPO REQUIREMENTS

Empirical group	HD hours/week (mean, 2024)	EPO IU/week (mean, 2024)
18–29, Female	10.9	7,867
18–29, Male	10.7	7,667
30–44, Female	11.2	7,226
30–44, Male	10.6	7,250
45–59, Female	11.2	8,036
45–59, Male	11.2	6,825
60–74, Female	10.8	7,148
60–74, Male	10.9	6,784
≥75, Female	10.0	7,250
≥75, Male	11.2	8,211

B. Grouping alternatives: coverage, structure, and implications for forecasting comparability

Unsupervised grouping methods could leave part of the cohort unassigned to any cluster (e.g., DBSCAN noise), which reduced the share of patients represented in forecastable segments. Because planning required cohort-level totals, this coverage effect was explicitly tracked.

Table 5 summarizes effective coverage for each grouping strategy (analyzable patients / base cohort), together with the number of groups and internal cohesion indicators (silhouette where applicable). DBSCAN(A) exhibited the lowest coverage due to a larger set of noise-labeled patients, whereas PAM, HCLUST, and DBSCAN(B) retained higher coverage with fewer groups.

TABLE V
GROUPING COMPARISON: NUMBER OF GROUPS, ANALYZABLE PATIENTS, COVERAGE, AND SILHOUETTE

Approach	Groups	Base cohort	Analyzable patients	Cov.	Silhouette
Empirical (age × sex)	10	2,592	1,735	66.94 %	–
HCLUST (Ward)	2	2,592	1,765	68.09 %	0.423
PAM	2	2,592	1,765	68.09 %	0.633
DBSCAN_A	14	2,592	1,616	62.35 %	0.532
DBSCAN_B	2	2,592	1,760	67.90 %	0.634

For downstream forecasting, lower coverage implied that segment-level forecasts would understate total demand unless excluded/noise patients are re-incorporated.

An add-back reconciliation step was therefore applied (Section II.F) so that all 2025 projections were comparable and represented the full cohort; adjustment magnitudes were reported with the 2025 outputs (Table 7). Silhouette values indicate moderate cluster cohesion, suggesting that grouping structures are not strongly separated, which may partially explain their lower forecasting performance relative to empirical stratification.

C. Holdout-year benchmarking (2024) and planning-oriented 2025 outputs

Forecasting configurations were screened using 2024 as a single-year holdout, and performance was assessed using weighted MAE and weighted sMAPE.

Table 6 reports the out-of-sample performance (wMAE) of the forecasting configuration retained for cross-grouping comparison. A weighted moving average (WMA) was used consistently across grouping strategies to preserve methodological comparability across outcomes.

TABLE VI
HOLDOUT (2024): BEST WITHIN-GROUPING PERFORMANCE (WMA) BY OUTCOME

Grouping	Patients (wMAE; w-sMAPE)	EPO IU/week (wMAE; w-sMAPE)	HD hours/week (wMAE; w-sMAPE)
Empirical	10.79; 18.7%	116,938.18; 25.0%	126.47; 18.7%
PAM	28.69; 13.2%	392,133.27; 22.6%	333.85; 14.3%
HCLUST	38.78; 12.0%	404,169.41; 18.7%	564.71; 12.9%
DBSCAN(A)	25.14; 53.5%	280,706.69; 60.2%	303.05; 49.7%
DBSCAN(B)	23.26; 10.7%	367,304.86; 21.3%	242.70; 10.6%

Under the fixed holdout protocol, the empirical stratification + WMA combination yielded the lowest error for all three outcomes, supporting its role as a robust, easy-to-implement baseline for short-horizon planning.

2025 adjusted projections and planning band

Because some grouping schemes excluded part of the cohort (rule-based exclusions in empirical; noise labeling in DBSCAN), an adjustment ("add-back") step was applied so that 2025 projections consistently represented the full cohort. Table 7 reports adjustment magnitudes and adjusted 2025 projections (Pred_2025*).

TABLE VII
2025 ADJUSTED PROJECTIONS (PRED_2025): PATIENTS, EPO IU/WEEK, AND HD HOURS/WEEK

Algorithm	Patient add-back	Pred_2025* Patients	Pred_2025* EPO (IU/week)	Pred_2025* HD hours/week
Empirical	30	608	4,991,668	6,882
PAM	–	555	4,429,902	6,444
HCLUST	–	598	4,592,067	6,621
DBSCAN(A)	149	740	5,946,197	8,115
DBSCAN(B)	5	558	4,459,948	6,517

For context, the work also reports a non-segmented base-model projection for 2025 (Table 8), which is systematically lower for all three outcomes in this dataset.

TABLE VIII
BASE-MODEL TOTAL PROJECTION FOR 2025 (PRED_2025)

Outcome	Pred_2025*
Patients	500
EPO (IU/week)	3,632,264
HD hours/week	5,767

To translate the dispersion observed across scenarios into a planning-oriented output, Table 9 operationalizes this variability by defining a recommended demand band. Specifically, the band is constructed by taking, for each planning period and indicator, the minimum and maximum values obtained among the adjusted scenarios, thereby producing an interval that can be directly used for capacity and resource planning. Under this convention, the lower bound corresponds to the PAM-adjusted scenario (minimum), while the upper bound is determined by the DBSCAN(A)-adjusted scenario (maximum).

TABLE IX
PLANNING BAND FOR 2025 (MIN – MAX ACROSS ADJUSTED SCENARIOS)

Band	Algorithm	Patients	EPO (IU/week)	HD hours/week
Minimum	PAM	555	4,429,902	6,444
Maximum	DBSCAN(A)	740	5,946,197	8,115

Figure 2 summarizes the adjusted 2025 scenario spread (Patients, EPO IU/week, HD hours/week), matching the planning-band logic in Table 9.

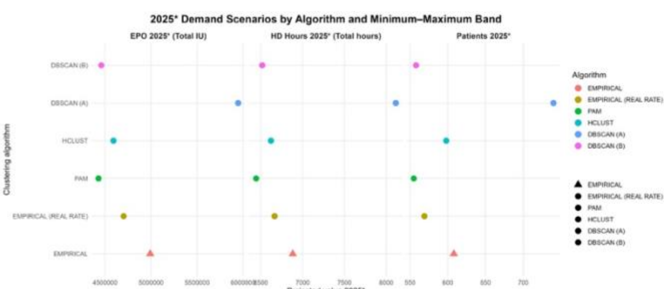

Fig 2 Adjusted 2025 demand scenarios by grouping strategy: projected patients, weekly EPO requirements (IU/week), and weekly HD hours (hours/week), with the implied min–max planning band.

IV. DISCUSSION

This work develops a reproducible workflow that maps routinely collected nephrology service records to planning-oriented demand indicators—patients/year and annual snapshots of cohort-level weekly requirements for HD hours and EPO IU—and evaluates whether patient grouping improves short-horizon forecasting under limited historical depth. Under the single-year holdout design (2024), empirical stratification (age × sex) paired with a parsimonious recency-weighted rule (WMA in exponential form) achieved the lowest weighted errors across all three outcomes (Table 6). This suggests that, in settings with sparse covariates and short annual histories, stable and interpretable segments can contribute more to planning accuracy than algorithmically derived partitions whose membership may be sensitive to administrative noise.

A central implication is that grouping strategies must be assessed not only in terms of feature-space structure but also in terms of effective coverage (Table 5). Density-based clustering can label atypical profiles as noise, and empirical rules can exclude patients due to missing operational fields; in both cases, segment forecasts aggregated without correction would mechanically understate total demand. The add-back reconciliation step (Section II.F) therefore serves as a comparability mechanism that restores cohort-level totals across grouping alternatives, and reporting adjustment magnitudes (Table 7) improves traceability of scenario differences.

Internal clustering indices (e.g., silhouette) are informative descriptively, but they do not guarantee temporal stability of aggregated demand within segments. In this application, the target series consists of yearly observations of cohort-level weekly requirements rather than dense weekly time series, which favors parsimonious smoothing rules once the segmentation yields sufficiently regular trajectories. This motivates focusing the main comparison table on a consistent forecasting rule across groupings: it isolates the effect of grouping on forecastability and avoids conflating segmentation differences with extensive model-selection variability.

Operationally, HD hours/week reflects prescription-driven machine and staffing requirements [7], while EPO demand represents a major recurring resource stream in anemia management with direct budget implications [14], [15]. The 2025 scenario dispersion and the resulting planning band (Table 9) should be interpreted as a transparent envelope induced by feasible grouping–forecasting configurations and reconciliation assumptions, rather than as a calibrated probabilistic prediction interval, providing a practical basis for procurement and capacity discussions.

Key limitations stem from the data and evaluation design: (i) a single-year holdout is a screening heuristic under short histories; (ii) missingness and analyzable definitions may bias segment intensities; (iii) deterministic add-back assumes excluded/noise patients share average intensity patterns, which may under- or over-estimate demand if these patients exhibit systematically higher or lower resource requirements; and (iv) the univariate scope does not explicitly capture structural breaks or exogenous drivers. Extensions include stability-aware grouping objectives aligned with forecasting performance, formal hierarchical reconciliation to enforce segment–total coherence, and decision-centric evaluation metrics linked to procurement risk and capacity shortfalls.

V. CONCLUSION

This study presents a reproducible decision-support workflow to anticipate short-horizon demand in public hemodialysis services using routinely collected clinical-administrative records from Paraguay. The analysis consolidates three planning indicators derived from the registry—yearly patient volumes (unique patients per year) and annual snapshots of cohort-level weekly requirements for hemodialysis workload (HD hours/week) and erythropoietin (EPO IU/week)—and reports them in a transparent, planning-oriented form that can support operational benchmarking and resource-planning discussions.

Methodologically, the work links (i) feasible patient grouping strategies that are compatible with routine information systems, (ii) a single-year holdout benchmark (2024) used as a consistent configuration-screening protocol under very limited time points, and (iii) translation of selected configurations into planning outputs for the 2025 horizon. Under this fixed screening design, empirical stratification (age group × sex) paired with a parsimonious recency-weighted smoothing rule (WMA in exponential form) yielded the lowest weighted errors across patients, HD hours/week, and EPO IU/week, supporting its role as a robust baseline for short-horizon institutional planning.

The comparison with unsupervised grouping alternatives (PAM, hierarchical clustering, and DBSCAN) highlights how grouping logics affect effective cohort coverage and, consequently, the comparability of aggregated demand projections. To ensure that all 2025 scenarios represent the full cohort, the workflow enforces coherent aggregation through an explicit add-back reconciliation step for excluded/noise patients, and reports the resulting adjusted projections. Interpreting the adjusted scenario dispersion as a planning envelope provides a practical output for procurement and capacity discussions, while preserving traceability to modeling choices and segment-level drivers.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the Instituto Nacional de Nefrología (INN) for providing access to the routinely collected clinical-administrative records used in this study and for the technical support that enabled data extraction and interpretation. The authors also thank the Ministerio de Salud Pública y Bienestar Social (MSPBS) for its institutional support and for facilitating the conditions required to conduct this work aimed at strengthening evidence-based resource planning in public hemodialysis services.

REFERENCES

[1] P. K.-T. Li, G. Garcia-Garcia, S. F. Lui, et al., "Kidney health for everyone everywhere—from prevention to detection and equitable access to care," *Kidney International*, vol. 97, no. 2, pp. 226–232, 2020, doi: 10.1016/j.kint.2019.12.002.

[2] National Kidney Foundation, "Global facts: About kidney disease," 2015.

[3] B. Bikbov, N. Perico, A. Remuzzi, et al., "Global, regional, and national burden of chronic kidney disease, 1990–2017: a systematic analysis for the Global Burden of Disease Study 2017," *The Lancet*, vol. 395, no. 10225, pp. 709–733, 2020, doi: 10.1016/S0140-6736(20)30045-3.

[4] M. C. González-Bedat, G. Rosa-Diez, and R. Correa-Rotter, "Advances in Hemodialysis in the Last Decade in Latin America," *Revista de Investigación Clínica*, vol. 75, no. 6, pp. 300–308, 2023, doi: 10.24875/RIC.23000224.

[5] Á. L. Martín de Francisco and V. Lorenzo Sellarés, "Enfermedad renal crónica," *Nefrología al Día* (Sociedad Española de Nefrología), 2024.

[6] National Institute of Diabetes and Digestive and Kidney Diseases, "Managing chronic kidney disease," 2016.

[7] W. Kim, "Hemodialysis procedure and prescription," in *The Essentials of Clinical Dialysis*, Y.-L. Kim and H. Kawanishi, Eds. Singapore: Springer, 2018, pp. 49–71.

[8] National Kidney Foundation, "KDOQI Clinical Practice Guidelines for Hemodialysis Adequacy," 2006.

[9] "In-center hemodialysis six times per week versus three times per week," *The New England Journal of Medicine*, vol. 363, no. 24, pp. 2287–2300, 2010, doi: 10.1056/NEJMoa1001593.

[10] J. T. Daugirdas, P. G. Blake, and T. S. Ing, *Handbook of Dialysis*, 5th ed. Philadelphia, PA, USA: Wolters Kluwer, 2017.

[11] European Medicines Agency, "Guideline on the quality of water for pharmaceutical use," EMA/CHMP/CVMP/QWP/496873/2018, 2020.

[12] Y. M. Gomes, R. B. Campos, A. M. Sant'ana, R. D. P. Keller, and S. T. A. Cassini, "Physical-chemical and microbiological characterization of water destined to hemodialysis," *Ambiente & Água*, 2024.

[13] L. S. Vest, P. Patel, and J. B. Patel, "Epoetin alfa," in *StatPearls*. Treasure Island, FL, USA: StatPearls Publishing, 2025.

[14] S. Sheingold, D. Churchill, N. Muirhead, A. Laupacis, R. Labelle, and R. Goeree, "The impact of recombinant human erythropoietin on medical care costs for hemodialysis patients," *Social Science & Medicine*, vol. 34, no. 9, pp. 983–991, 1992.

[15] A. D. Coutinho, M. Mahendran, M. DerSarkissian, S. A. Kitchen, C. F. Bell, M. Muoneke, and A. Richards, "Hemoglobin stability impact on healthcare resource utilization and costs among dialysis-dependent patients with anemia of end-stage kidney disease," *BMC Nephrology*, vol. 26, no. 1, p. 466, 2025.

[16] Instituto Nacional de Nefrología (INN), "Informe anual del Instituto Nacional de Nefrología: Hemodiálisis crónica 2010–2019," Ministerio de Salud Pública y Bienestar Social (Paraguay), 2019.

[17] Ministerio de Salud Pública y Bienestar Social – Instituto Nacional de Nefrología (Paraguay), "Informe anual 2022 del Instituto Nacional de Nefrología," 2022.

[18] Sociedad Paraguaya de Nefrología, "Cobertura de pacientes en Terapia de Reemplazo Renal en Paraguay (2010–2022)," *Revista de la Sociedad Paraguaya de Nefrología*, vol. 2, no. 2, pp. 10–18, 2023.

[19] Sociedad Paraguaya de Nefrología, "Pacientes en diálisis crónica en Paraguay 2010–2024: Prevalencia, incidencia, etiología y mortalidad," *Revista de la Sociedad Paraguaya de Nefrología*, vol. 3, no. 1, pp. 25–36, 2024.

[20] J. Pereira-Rodríguez, L. Boada-Morales, D. G. Peñaranda-Florez, and Y. Torrado-Navarro, "Diálisis y hemodiálisis: Una revisión actual según la evidencia," 2017.