Urban Water Journal

ISSN: 1573-062X (Print) 1744-9006 (Online) Journal homepage: www.tandfonline.com/journals/nurw20

From hourly readings to daily leak flags: an approach to post-meter leakage detection and analysis in residential water consumption

Ricardo Cobacho & Francisco J. Arregui

URBAN WATER JOURNAL
https://doi.org/10.1080/1573062X.2026.2625409

Taylor & Francis
Taylor & Francis Group

RESEARCH ARTICLE

OPEN ACCESS

From hourly readings to daily leak flags: an approach to post-meter leakage detection and analysis in residential water consumption

Ricardo Cobacho and Francisco J. Arregui

Department Hydraulic & Environmental Engineering, Universitat Politècnica de València, València, Spain

ABSTRACT

This work presents a data-efficient procedure for distinguishing between leak and no-leak days using standard hourly smart meter readings. The approach relies on a fixed Minimum Night Flow window and classifies days based on the presence or absence of continuous registered volume during night hours. Applied to more than 21,000 users over one year, it is validated through the application of a Generalised Linear Mixed Model, which confirms that leak days exhibit systematically higher hourly volumes across the usage groups. The procedure also quantifies the hourly and daily impact of leakage and characterises its temporal organisation and influence on diurnal consumption patterns. Using only routinely available hourly data, it provides utilities with a practical framework for assessing post-meter leakage and enhancing the operational value of existing smart-meter infrastructures.

ARTICLE HISTORY
Received 30 July 2025
Accepted 27 January 2026

KEYWORDS
Residential water consumption; leakage; hourly readings; smart metering

1. Introduction and objectives

Understanding residential water consumption at the user level is becoming increasingly crucial for demand management, system operation, and designing effective water-saving strategies (Cominola, Giuliani, et al. 2018; Rizzoli et al. 2014). Beyond aggregate volumes, utilities and regulators require information on when and how water is used to support peak demand mitigation, leakage detection, and targeted interventions (Heydari, Cominola, and Stillwell 2022; Mazzoni et al. 2024). In this context, the deployment of smart water meters has enabled the collection of consumption data at finer temporal resolutions, opening up new opportunities for characterising user behaviour and extracting meaningful consumption patterns (Randall and Koech 2019). However, the type and reliability of information that can be derived from smart meter data strongly depend on the temporal resolution of the measurements, which constrains both the scope of achievable analyses and their applicability in real-world operational settings (Cominola, Giuliani, et al. 2018).

The temporal resolution of smart water meter data plays a central role in determining the level of detail at which residential water use can be analysed. Very high-resolution data, ranging from seconds to sub-minute intervals, enable the identification of individual usage events and end uses, often supported by signal processing techniques or supervised machine learning approaches (Heydari, Cominola, and Stillwell 2022; Marsili et al. 2024; Mazzoni, Blokker, et al. 2024). These approaches can capture short-duration and overlapping events that are typically masked at coarser resolutions, enabling detailed behavioural analyses and, in some cases, user-level leakage detection (Zese et al. 2021). However, their practical applicability is often limited by the need for specialised sensing technologies, extensive training data, model calibration, and large data volumes, as well as by the restricted availability of high-resolution data in routine deployments (Alassio et al. 2024; Cominola, Giuliani, et al. 2018).

At coarser temporal resolutions, and particularly at the hourly scale, smart meter data are increasingly available through large-scale deployments and routine utility operations. Previous studies have shown that hourly readings are sufficient to analyse residential demand, capture behavioural transitions, support forecasting and stochastic modelling, and perform user segmentation and clustering at scale (Cominola, Spang, et al. 2018; Leyli-Abadi et al. 2017; Rizzoli et al. 2014; Romano and Kapelan 2014). Hourly data have also been used to evaluate demand-side management programmes and to analyse trends, anomalies, and long-term consumption patterns in large AMI datasets, including issues related to data quality and robustness (Khaki and Mortazavi 2022; Obaid et al. 2024; Pesantez, Berglund, and Kaza 2020; Randall and Koech 2019; Wang et al. 2024). Collectively, these studies indicate that hourly readings provide a suitable compromise between analytical capability and operational feasibility,

CONTACT Ricardo Cobacho rcobacho@ita.upv.es

making them particularly appropriate for large-scale user-level analyses.

Despite this progress, existing studies have not yet proposed a robust and scalable procedure to distinguish between leak and no-leak days at the individual user level using only hourly smart meter readings (Alassio et al. 2024). Most available approaches either rely on sub-hourly data, focus on aggregated indicators, or require multi-day analysis windows that are sensitive to missing data, which remains an inherent issue in AMI environments. As a result, the literature still lacks a methodology capable of identifying daily leakage occurrences at scale, quantifying their impact on registered consumption, and characterising the associated behavioural patterns without relying on fine-grained measurements or highly complete datasets.

The present work addresses this gap by exploiting the diagnostic potential of the Minimum Night Flow (MNF) period within hourly consumption profiles. The first contribution is the development and validation of a daily classification procedure to distinguish leak and no-leak days for each user, based on a fixed nocturnal window that can be consistently applied across heterogeneous consumption behaviours. The second contribution is the quantification of the additional registered volume associated with leak days, expressed at both hourly and daily scales. The third contribution is the characterisation of the temporal and behavioural effects of leakage, including its intensity, duration, recurrence, and influence on the diurnal distribution of consumption.

These contributions position the study as a large-scale and data-efficient approach to post-meter leakage analysis, grounded, to the authors' knowledge, in the most extensive user-level analysis conducted to date using hourly readings for this purpose. By relying exclusively on standard hourly data and avoiding methods that require high-frequency flow reconstruction or multi-day data blocks, the proposed methodology provides a practical and scalable framework that aligns with the operational conditions typically found in AMI systems.

2. Materials and methods

The successive calculation steps followed to achieve the stated objectives are described below.

2.1. Study area and data preparation

The proposed methodology was applied to an urban water supply system located in a Mediterranean city in Spain, serving approximately 250,000 inhabitants. The utility operates an AMI system that automatically transmits one hourly reading per connected customer. The dataset used in this study initially comprises 25,950 user accounts equipped with hourly-reading meters, representing approximately one-fourth of all customers in the service area. These accounts were selected because they offer sufficient temporal resolution for the intended analysis and correspond predominantly to standard urban households with relatively simple consumption profiles, consistent with the study's residential focus. Seasonal tourism has a minimal influence on this dataset, as tourist accommodations with substantial water use (e.g. hotels or rental complexes) generally fall outside the residential user segment considered here.

In this study, the registered volume refers to the hourly value measured by the meter. This quantity represents the volume registered above the meter's start-flow threshold; volumes associated with flows below that threshold are not detectable. In residential urban contexts, this limitation is minor: registered volume typically provides an accurate representation of actual consumption. For consistency, the term registered volume is therefore used throughout the rest of this paper. In the following, internal leakage refers to continuous, unintended flows that occur downstream of the customer meter and are recorded as registered volume. These flows usually originate from defective fixtures (e.g. toilet cisterns or taps) and may occasionally result from inadvertent continuous uses left open by users. Leakage in the distribution network upstream of the meter, as well as internal flows below the meter's start-flow threshold, fall outside the scope of this study. These exclusions reflect the operational capabilities of standard residential meters and do not restrict the objective of providing utilities with a practical and data-efficient way to extract leakage information from hourly registered volumes.

As is common in large-scale AMI deployments, the raw dataset exhibited anomalous readings due to intermittent transmission gaps or other failures. Such anomalous readings were defined as values that were blank, negative, or abnormally high. To identify unusually high values, a conservative statistical criterion based on the interquartile range (IQR) was applied separately for each hour of the day and for each month. For a given month and hour, an hourly reading was classified as anomalous if it exceeded the third quartile by more than three times the interquartile range. Data cleaning was performed following a sequential procedure applied independently to each user:

(1) Anomalous hourly readings (blank, negative, or abnormally high values) were identified for each day.

(2) When anomalies occurred in isolation (a single anomalous hourly value within an otherwise consistent day), the affected reading was corrected using the average value of the same hour computed over the remaining valid days of that week.
(3) When anomalies appeared in greater numbers on the same day, the entire day was discarded.
(4) The number of remaining valid days was evaluated for each user, and users with fewer than 300 valid full days were excluded from the analysis to ensure seasonal consistency.

In practice, the dominant causes of day invalidation were missing hourly readings (notably missing values at hour 0) and negative readings. These issues affected approximately 10,000 users at least once during the year. In contrast, extreme outliers were comparatively rare and limited to a few dozen users. Importantly, the presence of occasional anomalies did not necessarily lead to user exclusion. User removal was primarily driven by the accumulation of invalid days over time, in accordance with the validity criterion defined in point 4.

After this filtering stage, 4105 users were discarded, and the final dataset was reduced to 21,845 users, comprising nearly 7.7 million valid day records over the year. To account for behavioural differences, users were subsequently segmented according to their average daily registered volume. This data-driven grouping enables consistent application of the methodology across a heterogeneous user population without requiring socio-demographic information.

2.2. Characterisation of user groups

The analysis requires grouping users according to their overall usage level. For each user i, the average daily metered volume (L/d) is computed as:

$$\overline{C_i^{daily}} = \frac{1}{D_i} \sum\nolimits_{j \in D_i} \left(\sum\nolimits_{h=0}^{23} C_{ijh} \right) \quad (1)$$

where:

C_{ijh}: Registered volume (L/h) of user i on day j - during hour h

D_i: Set of valid daily records for user i

This quantity provides a stable indicator of long-term water use, supporting a consistent segmentation of the population. Users are subsequently assigned to groups defined by intervals of $\overline{C_i^{daily}}$, ensuring that structural differences in behaviour are adequately captured and that results remain comparable across usage contexts.

For each user group g, the mean hourly registered volume is calculated by averaging all valid day records within that group. For hour h, the group-level hourly mean (L/h) is:

$$\overline{C_{gh}^{hourly}} = \frac{1}{N_g} \sum\nolimits_{i \in g} \sum\nolimits_{j \in D_{gi}} C_{ijh} \quad (2)$$

where:

i: Each user belonging to the user group g

j: Each day belonging to D_{gi}

D_{gi}: Set of valid daily records for user i, from user group g

N_g: Number of valid record-days for user group g

To describe the shape of the daily pattern independently of magnitude, an hourly modulation coefficient is computed for each group:

$$M_{gh}^{hourly} = \frac{\overline{C_{gh}^{hourly}}}{\frac{1}{24 \cdot N_g} \sum\nolimits_{i \in g} \sum\nolimits_{j \in D_{gi}} \left(\sum\nolimits_{h=0}^{23} C_{ijh} \right)} \quad (3)$$

This coefficient normalises the hourly registered volume by the group-level average over the whole day. It is used to identify the hours with systematic minimal use, which later support the definition of the Minimum Night Flow window.

2.3. Identification of leak and no-leak days and periods

The identification of leak and no-leak days relies on defining an MNF window based on the hourly modulation coefficients obtained in Characterisation of user groups. For each application case, these coefficients are examined to determine the consecutive hours during which the registered volume systematically reaches its daily minimum, typically during nighttime. This interval, representing the period of consistently minimal consumption, is then adopted as the MNF window for the subsequent classification procedure. This MNF window may or may not be the same across user groups, depending on whether their temporal modulation patterns remain sufficiently consistent during the hours of minimal nighttime consumption, as happens in the case study analysed later. If such consistency does not hold and user groups exhibit markedly different patterns, the same procedure can still be applied by defining separate MNF windows for each group.

Each full day is labelled as a no-leak day if at least one of the MNF hours records a zero reading. Conversely, if all MNF readings are strictly positive, the day is classified as a leak day. Once days have been classified, consecutive sequences of the same category naturally form leak periods and no-leak periods, which are later analysed to characterise the temporal organisation of leakage.

The robustness of this classification requires validation along two complementary lines. First, a sensitivity analysis evaluates how the classification would vary if the criterion were relaxed, for instance, by allowing a small number of zero readings within the MNF window. Second, the proposed procedure is compared with existing approaches based on multi-day sliding windows, such as those used by Romano and Kapelan (2014) and Alassio et al. (2024). In these approaches, a leak period is typically identified only after at least 72 consecutive hours of continuous consumption have elapsed. It follows a similar logic but introduces additional complexity, as resulting leak periods do not generally align with full 24-hour days.

To perform this statistical validation, the mean hourly registered volume is modelled as a function of a leak indicator using a Generalised Linear Mixed Model (GLMM), fitted separately for each usage group (lme4 package in R):

$$\overline{C_{gij}^{hourly}} = \beta 0_g + \beta 1_g \cdot Leak_{ij} + u_{gi} + \varepsilon_{gij} \tag{4}$$

where:

$\overline{C_{gij}^{hourly}}$ (dependent variable): hourly average registered volume (L/h) of user i in group g on day j.

$\beta 0_g$ (fixed intercept): common mean hourly registered volume (L/h) for group g in no-leak days

$\beta 1_g$ (fixed coefficient of the variable $Leak_{ij}$): increase in mean hourly registered volume (L/h) associated with leak days in group g

$Leak_{ij}$ (independent variable): value that distinguishes no-leak days (0) from leak days (1) for each day j of each user i

u_{gi} (random effect): coefficient of each user i from group g representing the specific deviation (L/h) of that user from the group's $\beta 0_g$

ε_{gij} (residual error): variability (L/h) not explained by fixed or random effects

2.4. Quantification of the impact of leakage on registered volume

In addition to the statistical validation shown in the previous section, the GLMM provides a direct estimate of the contribution of leakage to the registered hourly volume. In the model specification, the intercept $\beta 0_g$ represents the average hourly registered volume of users in group g during no-leak days, while the coefficient $\beta 1_g$ quantifies the systematic difference in hourly registered volume between leak and no-leak days. Since the leak indicator is the only fixed effect, $\beta 1_g$ isolates the component of registered volume attributable to continuous background flow.

Leak magnitude can be expressed at two temporal scales. At the hourly scale, $\beta 1_g$ provides a direct estimate of the persistent flow associated with leakage (L/h). At the daily scale, the corresponding additional registered volume is obtained by multiplying this hourly value by 24, yielding a group-level estimate of the daily impact of leakage. These quantities offer a consistent basis for comparing leakage effects across users with different usage levels and behavioural patterns. User-specific deviations are accounted for by the random effects, ensuring that individual behavioural variability does not bias the estimation of leak magnitude. The resulting indicators constitute the quantitative reference for subsequent analyses of temporal organisation and usage pattern modification.

2.5. Analytical approach to characterising leakage influence

Once leak and no-leak days have been identified for every user, the analytical phase quantifies how internal leakage affects registered volume and daily consumption behaviour. This analysis integrates three complementary dimensions:

- First, the temporal organisation of leakage is examined by identifying the duration and recurrence of consecutive leak and no-leak periods. This allows characterising how leakage episodes emerge, persist, and disappear over time.
- Second, the influence of leakage on the diurnal modulation of registered volume is assessed by comparing hourly profiles derived from leak and no-leak days. This comparison highlights smoothing effects, shifts in typical usage patterns, and changes in the relative contribution of night and daytime consumption.
- The third dimension evaluates registered volume as a function of the average hourly flow rate, providing insight into how leakage alters the distribution of volumes across flow conditions and how persistent background flows reshape the lower end of the flow spectrum.

To quantify differences in the stability of hourly patterns, a variability analysis based on the Mean Squared Error (MSE) was carried out. For each user group and for each hour of the day, the MSE was computed using all valid daily observations in the year. Daily hourly values were compared against the corresponding hourly mean of the same day type (leak or no-leak), allowing the MSE to reflect the internal variability of each condition. Higher MSE values indicate greater dispersion of hourly consumption around its typical profile, whereas lower values denote more stable and consistent behaviour.

Together, these analyses provide a comprehensive framework for understanding the intensity, temporal structure, variability, and flow-related implications of internal leakage once the daily classification is established.

3. Results

3.1. Data availability and structure of the dataset

After the data-quality checks and corrections described above, a total of 21,845 users were retained for analysis. Table 1 summarises their distribution according to average daily water consumption and reports, for each group, the percentage of valid full days (i.e. days with all 24 hourly readings) relative to the 365 days of the year. This metric provides a direct indication of data completeness across user groups.

Groups 1 to 3 encompass the large majority of users and correspond to consumption ranges that are entirely consistent with typical residential water use. These groups, therefore, constitute the core of the analysis and the reference population for the residential patterns examined in this study. In contrast, groups 4 and 5 are characterised by substantially higher average daily consumptions and represent a very small fraction of the total user base. Although these groups were initially identified for completeness, they are likely to include non-residential or mixed-use connections that are not representative of residential water use, either in terms of consumption behaviour or population size. Consequently, the quantitative analyses presented in the remainder of the paper focus exclusively on groups 1–3, while results for groups 4 and 5 are not further discussed.

Table 1. Available days with 24 readings per user group.

Users group	Reg. Vol. (L/d)	# Users	% Users	% Days with all 24 readings
1	<100	6105	28%	97%
2	100–299	11,431	52%	96%
3	300–599	3825	18%	96%
4	600–999	352	2%	97%
5	≥1000	132	1%	96%
		21,845	100%	96%

The hourly modulation coefficients and the percentage of zero readings were computed considering the complete set of users, without distinction by consumption group (Figure 1). A contiguous nocturnal interval between 01:00 and 07:00 is identified as the period of systematically minimal and temporally uniform consumption. Although the modulation coefficients at the boundaries of this interval (hours 1 and 7) slightly exceed 0.5, they remain close to this value and form a narrow, stable range that is clearly separated from adjacent hours by marked increases in modulation (hours 0 and 8). This structure supports its definition as a Minimum Night Flow (MNF) window. While a shorter window would yield lower coefficients, adopting a longer nocturnal interval increases the robustness of the daily classification by requiring persistence of registered volume over more consecutive hours. For the case study analysed, the 01:00–07:00 interval is therefore adopted as the MNF window for subsequent analyses.

3.2. Daily leak/no-leak classification

Applying the MNF criterion, each valid day is classified as either a leak or a no-leak day. The resulting distribution across the relevant user groups (1-3), reported in Table 2, shows a low proportion of leak days (ranging from 2% to 14%), indicating that continuous unintended registered volume is uncommon in these predominantly residential users. Accordingly, Table 2 provides a descriptive summary

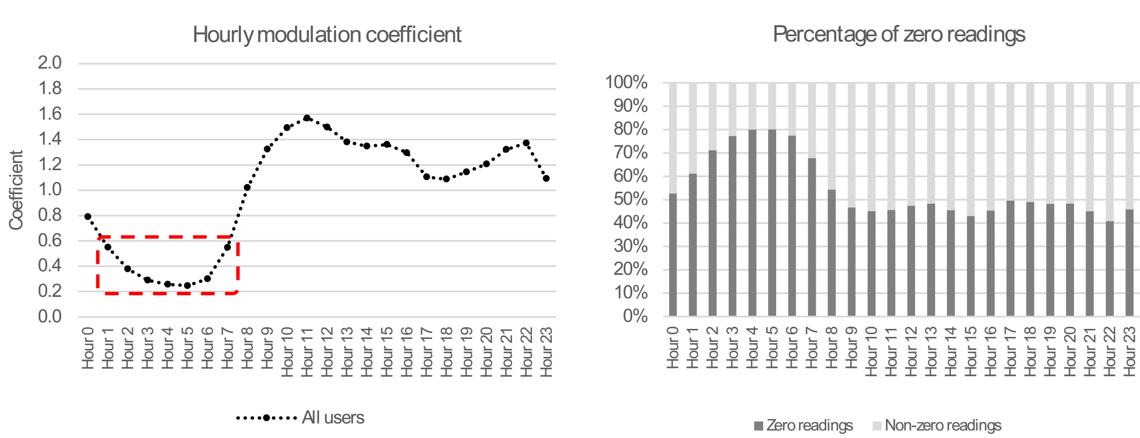

Figure 1. Average hourly modulation coefficient (left) and percentage of zero readings per hour (right) computed for the complete set of users.

Table 2. Leak and no-leak days (groups 1 to 3)

Users group	Leak days	No-leak days	Total days	% Leak	% No-leak
1	53,171	2,097,911	2,151,082	2%	98%
2	262,860	3,758,768	4,021,628	7%	93%
3	184,525	1,161,832	1,346,357	14%	86%

Table 3. Percentage of users, per group, by leak occurrence and rate of leak days (groups 1 to 3)

Users group	% Users (total)	% Users with no leaks	% Users with leaks	% Users with leaks (#users per % of leak days)									
				0–10%	10–20%	20–30%	30–40%	40–50%	50–60%	60–70%	70–80%	80–90%	90–100%
1	28%	60%	40%	34%	3%	1%	0%	1%	0%	0%	0%	0%	0%
2	52%	35%	65%	50%	6%	3%	2%	1%	1%	1%	1%	0%	1%
3	18%	22%	78%	51%	7%	4%	3%	3%	2%	2%	1%	1%	3%

Table 4. Percentage of days classified by the number of zero readings during night-hours (groups 1 to 3)

Users group	% Number of days with n night hours of zero readings						
	n = 1	n = 2	n = 3	n = 4	n = 5	n = 6	n = 7
1	1%	1%	1%	3%	8%	18%	69%
2	2%	3%	6%	13%	23%	28%	26%
3	4%	7%	13%	20%	24%	20%	12%

of the classification outcomes rather than a performance assessment against an external ground truth.

The classification can also be evaluated at the user level. Table 3 displays the percentage of users falling into each leak-rate category (0–10%, 10–20%. . .). Most users experience leaks only occasionally, although the proportion increases with the average registered-volume group.

To assess the robustness of the classification criterion, Table 4 reports the percentage of days containing 1 to 7 zero readings during the MNF hours for the user groups 1 to 3. Days with only one zero reading are infrequent (typically between 1% and 4%), indicating that the potential impact of misclassification due to an isolated zero value is negligible.

3.3. Statistical validation of the classification

Table 5 shows that the statistical validation is strongly supported across the analysed user groups (1–3), with large t-values for the leak coefficient (t-Leak), thus confirming that the distinction between leak and no-leak days is highly significant. The coefficient $\beta 1$ is consistently positive and large relative to its standard error, indicating that days classified as leak present systematically higher hourly registered volume than no-leak days. These effects are observed in both nighttime and daytime registered volumes, demonstrating that the classification method performs robustly across different residential consumption levels and behavioural patterns.

3.4. Magnitude and impact of leakage

Table 6 provides the estimated registered volume on no-leak days ($\beta 0$) and the additional flow attributable to leakage ($\beta 1$) for the user groups analysed (1–3). The increase associated with leak days is clearly identifiable across all three groups, with typical values ranging from approximately 5 to 8 L/h, consistent with modest but persistent background flows.

Table 5. Results of statistical validation of leak and no-leak days according to GLMM (groups 1 to 3)

Users group	Time	FIXED EFFECTS			
		β0 - Intercept (±Error)	t Intercept	β1 Leak (±Error)	t Leak
1	Nighttime	0.41 (±0.01)	46.7	6.19 (±0.02)	331.2
1	Daytime	2.23 (±0.02)	96.2	4.71 (±0.02)	224.3
1	Full day	1.70 (±0.02)	94.9	5.17 (±0.02)	293.2
2	Nighttime	1.90 (±0.02)	117.8	7.50 (±0.02)	429.6
2	Daytime	9.77 (±0.03)	337.5	5.20 (±0.02)	265.2
2	Full day	7.47 (±0.02)	336.3	5.88 (±0.02)	356.9
3	Nighttime	4.08 (±0.06)	67.8	9.95 (±0.04)	256.1
3	Daytime	19.73 (±0.07)	281.2	6.99 (±0.04)	170.8
3	Full day	15.17 (±0.06)	275.3	7.84 (±0.03)	225.0

Table 6. Estimates for base registered volume and leakage impact according to GLMM analysis (groups 1 to 3)

	Base Reg. Volume (No-leak days)			Increase due to leak occurrence (Leak days)			
User group	L/h	L/day	% No-leak days	L/h	L/day	%	% Leak days
1	1.7	41	98%	5.2	125	306%	2%
2	7.5	180	93%	5.9	142	79%	7%
3	15.2	365	86%	7.8	187	51%	14%

Table 7. Average hourly registered volume during night- and day-hours (groups 1 to 3)

	Average hourly registered volume (L/h)			
	Night-hours (01:00 - 07:00)		Day-hours (08:00 - 00:00)	
User group	Leak days	No-leak days	Leak days	No-leak days
1	5.64	0.43	6.13	2.25
2	8.12	1.99	13.17	9.89
3	12.19	4.37	23.62	20.22

Table 7 summarises the average hourly registered volume for leak and no-leak days during nighttime hours (1–7 AM) and daytime hours (8–0 AM) for all user groups. In the residential groups (1–3), leak days systematically show higher hourly flows than no-leak days, with the contrast being much sharper at night than during the day, confirming that the MNF window concentrates the leakage signal.

3.5. Duration and frequency of leak and no-leak periods

For each user, consecutive leak and no-leak days were grouped into periods. Here, a 'period' denotes a contiguous sequence of days of the same type (leak or no-leak), so that leak and no-leak periods alternate by construction and do not represent independent physical leak events. Figure 2 (left) shows the average number of periods per user and per year. As expected, leak and no-leak periods appear in similar numbers within each group because both types alternate by construction. The relevant variation lies across groups: Group 3 shows a higher number of periods, indicating more frequent transitions between leak and no-leak states. Figure 2 (right) presents the mean durations, with leak periods consistently short (only a few days) and no-leak periods substantially longer, reflecting the intermittent nature of most leak events.

3.6. Effect of leakage on the daily shape and variability of hourly profiles

A detailed analysis of hourly patterns is conducted for the residential user groups analysed (1-3), which represent typical residential behaviour and account for the majority of the dataset. These groups provide the most consistent context for interpreting continuous unintended flows. The hourly heatmaps (Figure 3) summarise the 24-hour registered-volume pattern for leak and no-leak days in these residential groups. Leak days show systematically higher nighttime values and a smoother diurnal shape.

Table 8 quantifies these differences shown in Figure 3. All variability indicators (standard deviation, coefficient of variation, interquartile range, entropy and peak-to-valley range) are lower during leak days, indicating a dampened hourly modulation. The entropy metric shows a slight but consistent increase in leak days. This small increase in entropy is consistent with a more uniform distribution of registered volume across hours during leak days, despite their lower dispersion indicators. In all user groups except group 1, leak days also exhibit a reduced peak-valley range, reflecting the smoothing effect of continuous background flow. In group 1, whose consumption levels are extremely low and often close to zero on no-leak days, the peak-valley value can be slightly higher during leak days due to the minimal absolute differences involved. This behaviour is specific to the low-consumption nature of group 1 and does not affect the overall trend observed in the remaining residential groups.

Figure 4 shows the hourly MSE profiles for leak and no-leak days in groups 1-3. The separation between the two conditions is most apparent during the MNF hours, particularly in the higher-usage groups.

3.7. Effect of leakage on the distribution of registered volume across flow rates

Figure 5 shows the volume-flow distributions for night and day hours. Across the analysed user groups, leak

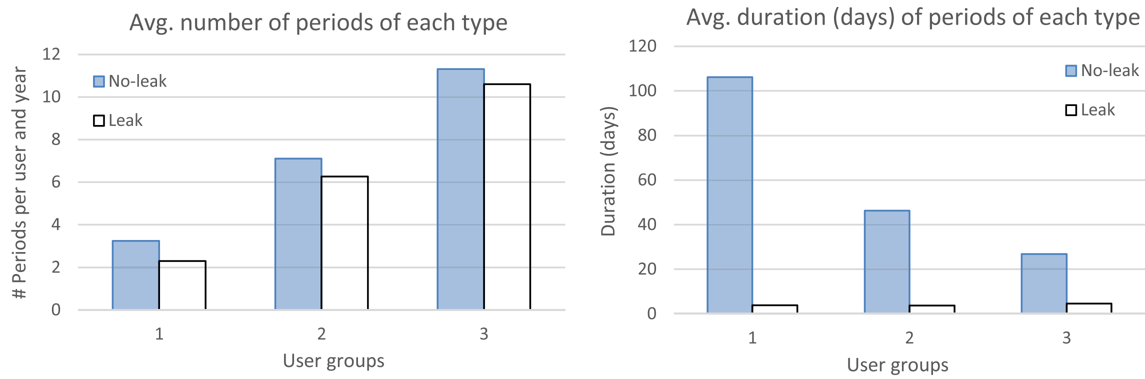

Figure 2. Average number (left) and average duration (right) of leak and no-leak periods (groups 1-3)

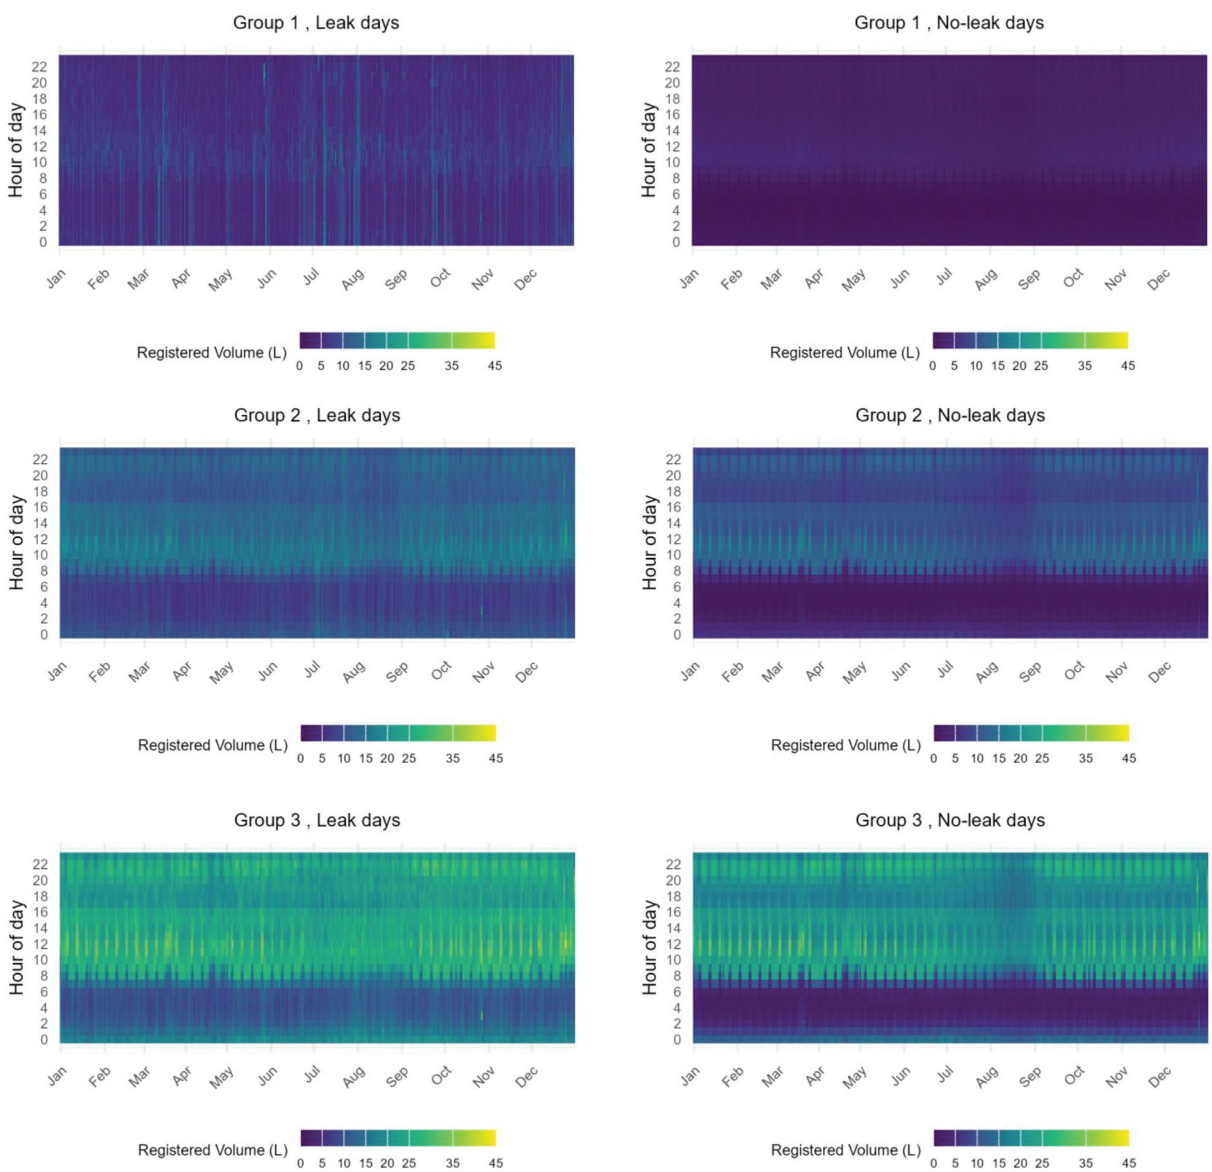

Figure 3. Heat maps of hourly registered volume for leak and no-leak days (groups 1 to 3)

Table 8. Statistical descriptors for the heatmaps (groups 1 to 3)

Users group	Day type	Mean SD	Mean CV	Mean IQR	Mean Entropy	Mean Peak-to-Valley
1	Leak	1.21	0.20	1.61	3.16	4.41
	No-leak	1.01	0.59	1.56	2.98	3.16
2	Leak	2.96	0.25	4.28	3.15	9.75
	No-leak	4.24	0.56	6.57	3.00	13.01
3	Leak	6.22	0.31	9.51	3.13	19.44
	No-leak	8.33	0.53	12.80	3.01	25.38

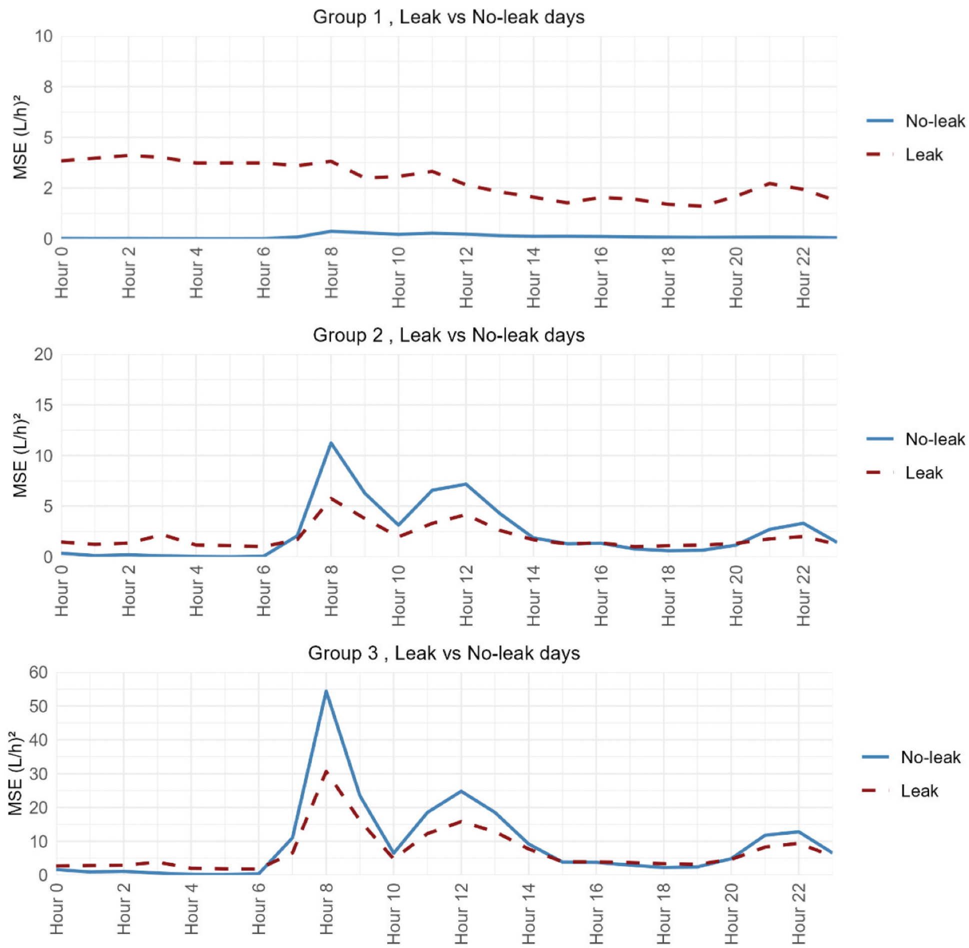

Figure 4. Mean hourly MSE values for leak and no-leak days(groups 1 to 3)

days shift the distribution toward higher average hourly flow rates and reduce the share of volumes concentrated at low flows, particularly during the MNF window.

Figure 6 presents the cumulative registered volume as a function of the average hourly flow rate for user groups 1 to 3. During night hours, the proportion of volume recorded at low flow rates is markedly smaller on leak days, reflecting the contribution of background flow. The three groups exhibit similar patterns, with curve compression increasing as the consumption level rises.

Figure 7 classifies the registered-volume days according to their average hourly flow during night and day hours. Leak days are concentrated in higher-flow categories, whereas no-leak days dominate the lower ranges, reinforcing the contrast between the two types of behaviour.

4. Discussion

The proposed daily classification method provides a structured and reliable framework for distinguishing between leak and no-leak days using hourly registered-volume data, revealing systematic quantitative differences between the two conditions at scale. In particular, leak days are consistently associated with higher

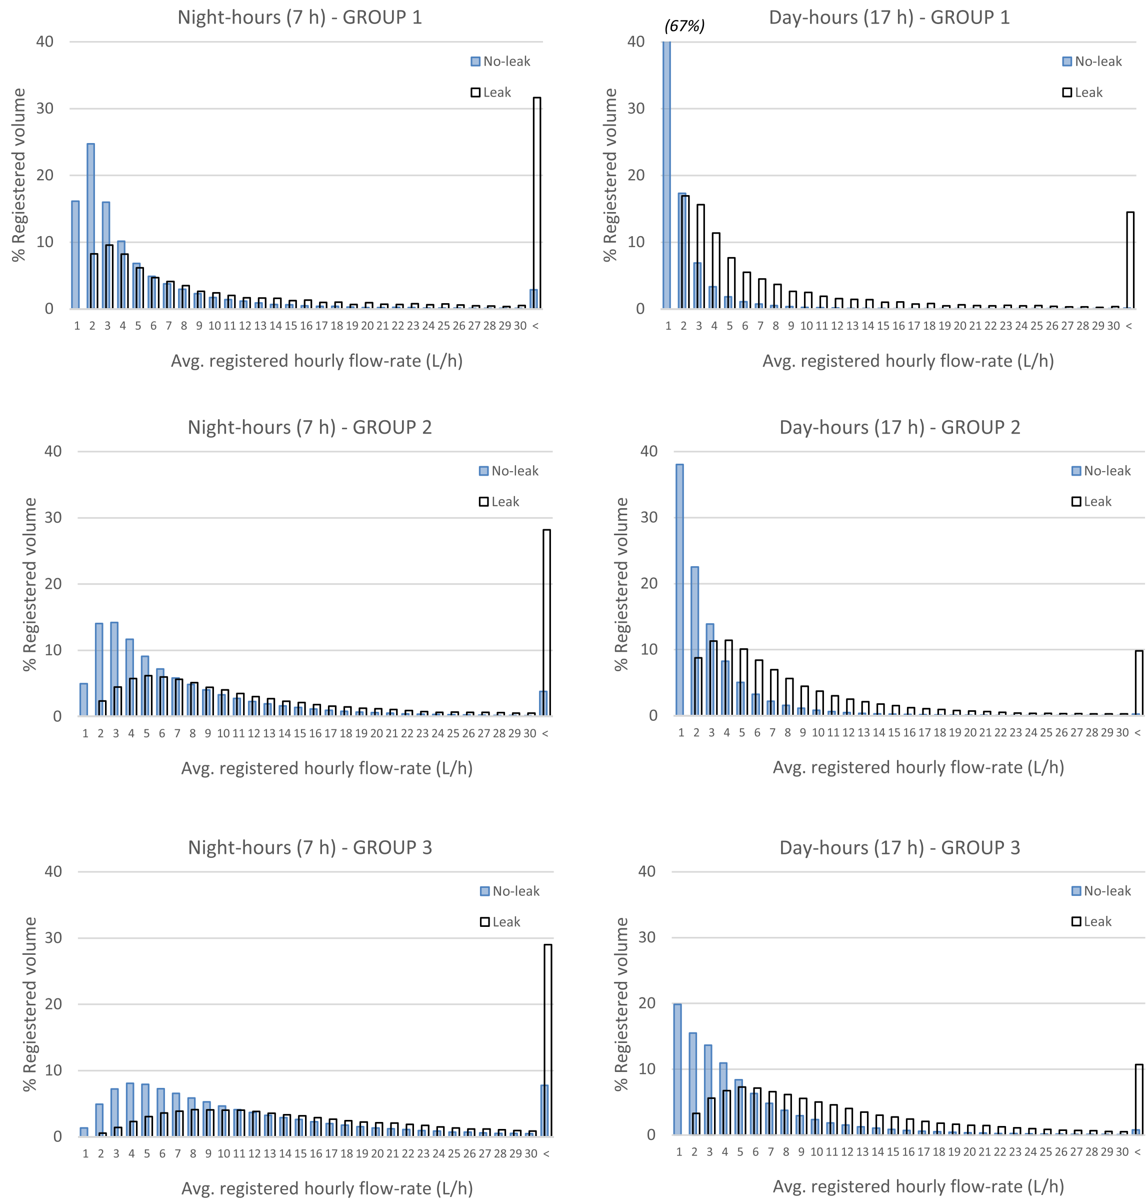

Figure 5. Volume-flow patterns for leak and no-leak days (groups 1 to 3)

registered volumes, with typical increases of 5–8 L/h in residential users, as confirmed by statistical validation. By basing the classification exclusively on the presence of continuous non-zero readings during the MNF period, the method leverages the most stable segment of the daily usage cycle to identify anomalous behaviour. The definition of leakage adopted here is necessarily constrained by what standard residential meters can record, that is, continuous unintended flows above the start-flow threshold; this operational definition supports the practical focus of the method, which aims at extracting reliable leakage signals from routinely available hourly data rather than capturing every possible form of internal loss. The results in Table 2 demonstrate that this criterion effectively distinguishes between two distinct types of days. The statistical validation provided by the GLMM in Table 7 confirms that the differences in hourly registered volume between the two categories are substantial and statistically significant. This framework demonstrates that meaningful leakage detection can be achieved with a short, fixed window, a contribution reinforced by the large dataset analysed, which supports the robustness and general applicability of the procedure.

The MNF window (01:00 - 07:00) can be adjusted to reflect local patterns of minimal use and, if needed, defined separately across user groups when nighttime behaviour differs. In all cases, the statistical validation

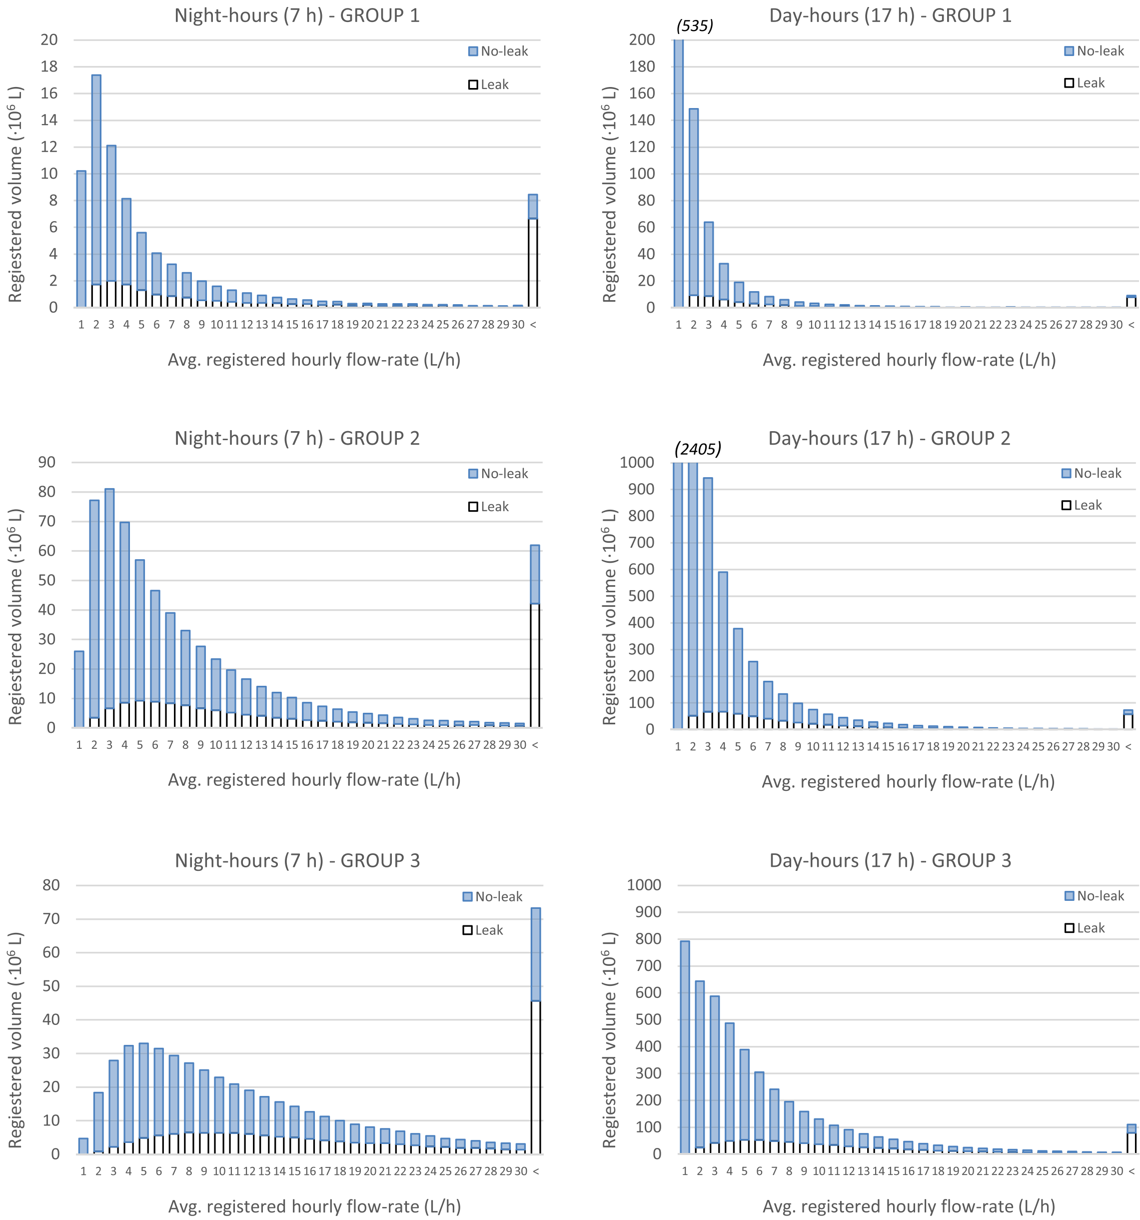

Figure 6. Cumulative registered volume over hourly flow rate (groups 1 to 3)

proposed here provides a consistent mechanism to assess whether such adaptations remain meaningful and robust.

However, the procedure has some inherent limitations that should be acknowledged. It cannot operate when readings are missing within the MNF window, since this period is structurally required for the daily classification and may reduce representativeness in systems with low nighttime transmission reliability. Moreover, the method detects continuous registered volume rather than physical leakage, and the hourly resolution restricts the identification of very small or intermittent events. These constraints arise from the temporal granularity and availability of the data rather than from the design of the proposed method.

The method also assumes that the MNF window corresponds to a period of minimal intentional use, which is appropriate for the vast majority of residential users. Households with scheduled nighttime consumption, such as automatic irrigation, fall outside the main scope of the study and correspond to minority consumption profiles. Such cases would also challenge alternative approaches based on multi-day sliding windows, since any prolonged intentional nighttime use produces continuous hourly flow. This limitation is therefore inherent to the resolution and nature of the available data rather than to the specific design of the proposed method.

Once the classification is established, the magnitudes of leakage effects can be interpreted across the analysed

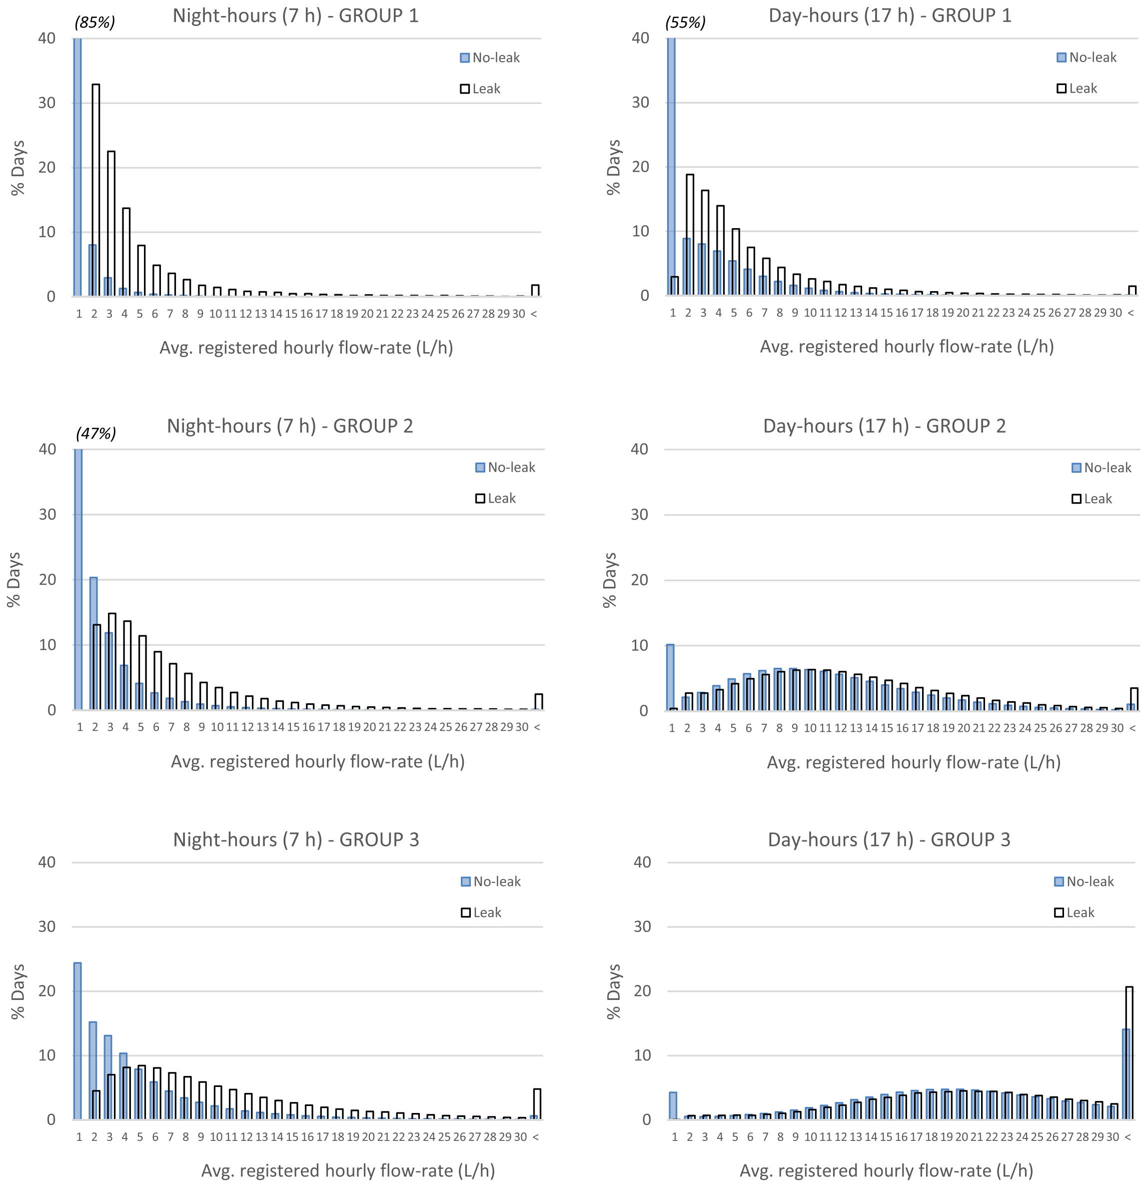

Figure 7. Daily registered volume patterns for leak and no-leak days (groups 1 to 3)

user groups. The GLMM coefficients in Table 5 indicate statistically significant increases in nighttime flow during leak days, amounting to approximately 6–10 L/h. These hourly effects translate into substantial differences in daily registered volumes. For instance, Table 6 shows an 80% increase in average daily volume between no-leak and leak days for group 2. The direction and magnitude of these differences align with the descriptive statistics presented in Table 8. Although variability increases with average consumption level, the leak signal remains robust, reflecting the suitability of the classification procedure across residential usage contexts.

Understanding the temporal organisation of leak and no-leak periods provides further insight into user behaviour. Figure 2 illustrates that residential users experience relatively few leak episodes per year, typically of short duration. The number of leak and no-leak periods shown in Figure 2 reflects the frequency of transitions between the two states, rather than the occurrence of independent physical leak events. The higher number of periods observed in group 3, therefore, indicates more frequent alternation between leak and no-leak days. This behaviour is consistent with intermediate-consumption households, where small background flows may intermittently appear or disappear due to minor fixture malfunctions or occasional nighttime uses above the meter start-flow threshold. As average consumption increases, non-zero nighttime readings become more frequent, making the classification more sensitive to such changes. Complementary analyses of hourly modulation patterns (Figures 3 and 4) reveal that leak days exhibit reduced variability throughout the day, with lower peak-to-valley amplitudes and smaller dispersion metrics across all statistical indicators reported in Table 8. This smoothing effect is consistent with the presence of a persistent background flow and is especially noticeable in the higher-usage groups.

Furthermore, the MSE analysis (Figure 4) reinforces the importance of the early-morning interval, as deviations between leak and no-leak conditions are most significant between 01:00 and 07:00. In the intermediate and higher residential groups (2 and 3), the higher daytime MSE observed on no-leak days reflects the greater influence of heterogeneous user behaviour. In contrast, the presence of a persistent background flow during leak days smooths hourly patterns despite higher mean consumption.

Taken together, Figures 5, Figures 6 and 7 reveal what can be described as a statistical signature of leakage in hourly registered-volume data. Rather than manifesting as isolated peaks or discrete events, leakage produces a systematic restructuring of the hourly flow distribution, primarily affecting its lower end. Across all residential groups, leak days are characterised by the suppression of very low or near-zero average hourly flows and by a consistent shift of registered volume toward persistent background flow rates. Figure 5 illustrates this effect through the redistribution of volume across flow classes, showing how leakage progressively replaces low-flow regimes, particularly during night hours. Figure 6 reinforces this interpretation by demonstrating that, on leak days, the cumulative registered volume accumulates more rapidly at moderate and high flow rates. In contrast, on no-leak days, a substantial share of the volume is concentrated at the lowest flows. Figure 7 complements this view at the daily scale, showing that days classified as leak systematically occupy higher average-flow categories, while no-leak days dominate the lowest ranges. Importantly, this signature is consistent across user groups despite differences in absolute consumption levels, indicating that the detection does not rely on arbitrary volume thresholds but on a stable structural pattern induced by continuous background flow. Within this framework, the absence of values in the 0–1 L/h class for leak days is a direct consequence of the classification itself: because leak days contain no zero readings during the MNF window, their averaged hourly flows cannot fall within this interval.

5. Conclusion

This study presents a scalable and data-efficient procedure for identifying post-meter leakage using standard hourly smart-meter readings. By classifying each day as leak or no-leak based on a fixed Minimum Night Flow window, the method enables robust detection without requiring sub-hourly data, complex event reconstruction, or multi-day analysis windows. Applied to more than 21,000 residential users over one year, the approach proves fully compatible with the data availability and reliability constraints of large-scale AMI systems.

The results show that leak days are consistently associated with modest but persistent background flows, typically in the range of 5–8 L/h, leading to substantial increases in daily registered volume. Leakage events are generally intermittent and short-lived, which helps explain their low detectability by users and highlights the value of automated monitoring. The statistical validation confirms that the daily classification captures meaningful and systematic differences in consumption behaviour across residential usage groups.

From an operational perspective, the proposed method provides utilities with a practical tool to extract leakage information directly from routinely collected hourly data, independent of manufacturer-specific meter alarms. It supports improved demand analysis, customer profiling, and asset management, while maximising the value of existing smart-meter infrastructure. Overall, the study demonstrates that reliable and informative post-meter leakage analysis can be achieved at scale using only standard hourly consumption data.

Acknowledgements

The authors gratefully acknowledge the reviewers for their thorough and constructive reviews, which contributed to a more precise and focused presentation of the study.

All authors contributed equally to the conceptualisation, methodology, analysis, writing, and revision of this manuscript.

Author contributions

CRediT: **Ricardo Cobacho:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing; **Francisco J. Arregui:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data availability statement

The dataset used in this study is confidential and cannot be shared due to restrictions imposed by the data provider. However, additional complementary aggregated outputs have also been generated that, while not essential to the core findings reported here, further illustrate the patterns analysed. The authors would be pleased to share these materials or explore collaborative work based on the same dataset, subject to the limitations imposed by confidentiality and data protection regulations.

Ethical approval

This study did not involve human participants, animals, or personal data; therefore, no ethical approval was required.

References

Alassio, S., V. Marsili, F. Mazzoni, and S. Alvisi. 2024. “Exploring Residential Minimum Night Consumption in a Real Water Distribution Network Based on Smart-Meter Data.” *Discover Water* 4 (1): 88. https://doi.org/10.1007/s43832-024-00150-5.

Cominola, A., M. Giuliani, A. Castelletti, D. E. Rosenberg, and A. M. Abdallah. 2018. “Implications of Data Sampling Resolution on Water Use Simulation, End-Use Disaggregation, and Demand Management.” *Environmental Modelling and Software* 102:199–212. https://doi.org/10.1016/j.envsoft.2017.11.022.

Cominola, A., E. S. Spang, M. Giuliani, A. Castelletti, J. R. Lund, and F. J. Loge. 2018. “Segmentation Analysis of Residential Water-Electricity Demand for Customized Demand-Side Management Programs.” *Journal of Cleaner Production* 172:1607–1619. https://doi.org/10.1016/j.jclepro.2017.10.203.

Heydari, Z., A. Cominola, and A. S. Stillwell. 2022. “Is Smart Water Meter Temporal Resolution a Limiting Factor to Residential Water End-Use Classification? A Quantitative Experimental Analysis.” *Environmental Research: Infrastructure and Sustainability* 2 (4): 045004. https://doi.org/10.1088/2634-4505/ac8a6b.

Khaki, M., and N. Mortazavi. 2022. “Water Consumption Demand Pattern Analysis Using Uncertain Smart Water Meter Data.” *International Conference on Agents and Artificial Intelligence* 3:436–443. https://doi.org/10.5220/0010834900003116.

Leyli-Abadi, M., A. Samé, L. Oukhellou, N. Cheifetz, P. Mandel, C. Féliers, and O. Chesneau. 2017. “Predictive Classification of Water Consumption Time Series Using Non-Homogeneous Markov Models.” *Proceedings - 2017 International Conference on Data Science and Advanced Analytics, DSAA 2017, 2018-January*: 323–331. https://doi.org/10.1109/DSAA.2017.32.

Marsili, V., F. Mazzoni, S. Alvisi, and M. Franchini. 2024. “From Pressure to Water Consumption: Exploiting High-Resolution Pressure Data to Investigate the End Uses of Water.” *Water Resources Management* 38 (13): 4969–4985. https://doi.org/10.1007/s11269-024-03898-6.

Mazzoni, F., M. Blokker, S. Alvisi, and M. Franchini. 2024. “An Enhanced Method for Automated End-Use Classification of Household Water Data.” *Journal of Hydroinformatics* 26 (2): 408–423. https://doi.org/10.2166/hydro.2024.125.

Mazzoni, F., V. Marsili, S. Alvisi, and M. Franchini. 2024. “Detection and Pre-Localization of Anomalous Consumption Events in Water Distribution Networks Through Automated, Pressure-Based Methodology.” *Water Resources and Industry* 31:31. https://doi.org/10.1016/j.wri.2024.100255.

Obaid, S., K. Hosoi, N. M. Ngoc, T. Inoue, and K. Yokota. 2024. “Assessment of Water Consumption Behavior in Single Households Using Smart Water Meters.” *Applied Sciences* 14 (19): 8857. https://doi.org/10.3390/app14198857.

Pesantez, J. E., E. Z. Berglund, and N. Kaza. 2020. “Smart Meters Data for Modeling and Forecasting Water Demand at the User-Level.” *Environmental Modelling and Software* 125:125. https://doi.org/10.1016/j.envsoft.2020.104633.

Randall, T., and R. Koech. 2019. “Smart Water Metering Technology for Water Management in Urban Areas.” *Water E-Journal* 4 (1): 1–14. https://doi.org/10.21139/wej.2019.001.

Rizzoli, A. E., A. Castelleti, A. Cominola, P. Fraternali, A. Diniz dos Santos, B. Storni, R. Wissmann-Alves, et al. 2014. *Presenter/Author Information Presenter/Author Information the SmartH2O Project and the Role of Social Computing in Promoting Efficient Residential Water Use: A First Analysis*. http://www.iemss.org/society/index.php/iemss-2014-proceedings.

Romano, M., and Z. Kapelan. 2014. “Adaptive Water Demand Forecasting for Near Real-Time Management of Smart Water Distribution Systems.” *Environmental Modelling and Software* 60:265–276. https://doi.org/10.1016/j.envsoft.2014.06.016.

Wang, R., X. Zhao, H. Qiu, X. Cheng, and X. Liu. 2024. “Uncovering Urban Water Consumption Patterns Through Time Series Clustering and Entropy Analysis.” *Water Research* 262:122085. https://doi.org/10.1016/j.watres.2024.122085.

Zese, R., E. Bellodi, C. Luciani, and S. Alvisi. 2021. “Neural Network Techniques for Detecting Intra-Domestic Water Leaks of Different Magnitude.” *IEEE Access* 9:126135–126147. https://doi.org/10.1109/ACCESS.2021.3111113.