

Inferring the Hidden: Privacy Risks of Microaggregation in Smart Meter Data

DEJAN RADOVANOVIC, Salzburg University of Applied Sciences and Paris Lodron University of Salzburg, Austria

JOAQUIN DELGADO FERNANDEZ, SnT - Interdisciplinary Centre for Security, Reliability and Trust, Luxembourg

MAXIMILIAN SCHIRL, Salzburg University of Applied Sciences, Austria

GUENTHER EIBL, Salzburg University of Applied Sciences, Austria

ANDREAS UNTERWEGER, Salzburg University of Applied Sciences and Paris Lodron University of Salzburg, Austria

SERGIO POTENCIANO MENCI, SnT - Interdisciplinary Centre for Security, Reliability and Trust, Luxembourg

Smart meter data, while essential for energy systems, pose significant privacy risks due to the behavioral information embedded in household electricity consumption patterns. Microaggregation has emerged as a promising anonymization technique to mitigate these risks. However, it remains unclear whether such aggregated profiles retain an identifiable structure that enables group membership inference while maintaining utility as it perturbs the data.

In this paper, we present a replicable methodology to evaluate the trade-off between utility and privacy in micro-aggregated smart meter data. We assess utility through household-level day-ahead load forecasting and evaluate privacy by implementing an unsupervised group membership inference attack. The attack combines distance-based record linkage with a two-stage majority voting scheme and is applied across a range of anonymity levels ($k = 5$ to 200) using both domain-specific features and deep neural representations.

Our results reveal a utility-privacy trade-off: while forecasting accuracy degrades only moderately (maximum 14% loss), group membership inference remains highly effective at lower k values, with success rates up to 80 times higher than random guessing.

These findings indicate that structural patterns persist through aggregation and can be exploited by adversaries, even without household-level identification, to enable targeted advertising, discriminatory profiling, or dynamic pricing. As such, microaggregation provides meaningful privacy protection at sufficiently higher k levels, underscoring the need for context-aware deployment in energy data sharing.

CCS Concepts: • **Security and privacy** → **Data anonymization and sanitization**; • **Computing methodologies** → **Unsupervised learning**; • **Information systems** → **Data analytics**.

Additional Key Words and Phrases: smart meter data anonymization, microaggregation, k-anonymity, utility-privacy trade-off, group membership inference, time series privacy, unsupervised learning

Availability of Data and Material:

The data and source code are available on request.

1 INTRODUCTION

Smart meters have become indispensable to modern energy infrastructures, enabling the continuous, high-resolution monitoring of household electricity consumption [69]. Their global deployment has accelerated, particularly across the European Union (EU) [29, 66]

Authors' addresses: Dejan Radovanovic, dejan.radovanovic@fh-salzburg.ac.at, Salzburg University of Applied Sciences and Paris Lodron University of Salzburg, Austria; Joaquin Delgado Fernandez, joaquin.delgadofernandez@uni.lu, SnT - Interdisciplinary Centre for Security, Reliability and Trust, Luxembourg, Luxembourg; Maximilian Schirl, maximilian.schirl@fh-salzburg.ac.at, Salzburg University of Applied Sciences, Austria; Guenther Eibl, guenther.eibl@fh-salzburg.ac.at, Salzburg University of Applied Sciences, Austria; Andreas Unterweger, andreas.unterweger@fh-salzburg.ac.at, Salzburg University of Applied Sciences and Paris Lodron University of Salzburg, Austria; Sergio Potenciano Menci, sergio.potenciano-menci@uni.lu, SnT - Interdisciplinary Centre for Security, Reliability and Trust, Luxembourg, Luxembourg.

and the United States [62], and in some jurisdictions, such as Luxembourg, smart meter coverage surpassed 99% as early as 2023 [15]. Typically, these devices record energy usage at intervals of 15 to 30 minutes [28], generating fine-grained time series data—referred to as **load profiles**—comprising timestamped power consumption values.

Beyond enabling real-time monitoring, smart meter data supports a range of applications, including grid optimization, enhanced energy efficiency, and predictive analytics [69]. However, the same granular detail that enables such functionality also introduces serious privacy risks. Load profiles can expose behavioral patterns such as occupancy [10, 36, 51], daily routines [1, 47, 49], and appliance usage [46, 48, 58]. Moreover, this data can be correlated with socio-demographic attributes—such as dwelling type, location [5, 41, 68], home ownership status [4, 6], or the presence of energy-intensive amenities like saunas and swimming pools [37, 56]. When combined with external datasets, these profiles can yield highly detailed household characteristics, substantially amplifying privacy risks.

Given these amplified risks, regulatory frameworks such as the European Union's General Data Protection Regulation (GDPR) classify smart meter data as personal data [33]. Accordingly, data controllers—such as energy suppliers or system operators—must apply Privacy Enhancing Technologies (PETs) before sharing smart meter data to reduce the risk of privacy threats like identifying individuals, **linking load profiles across datasets**, or inferring sensitive information [70].

Addressing these concerns demands a nuanced balance: preserving individual privacy while retaining sufficient data utility for essential energy management tasks. While techniques such as pseudonymization maintain high analytical utility, they often fall short in preserving privacy. Conversely, stronger protections like random data aggregation can obscure critical temporal details, limiting the usefulness of the data for fine-grained applications such as time-of-use pricing, demand response and load forecasting [20, 38]. Load forecasting is particularly relevant as it enables energy suppliers to anticipate demand fluctuations, stabilize grid operations, and optimize energy procurement in advance, thereby playing a key role in achieving cost-efficient and sustainable energy management [55].

At the same time, the EU is accelerating digitalization demands ever-larger datasets to fuel advanced AI-driven insights and data spaces [31]. Complementary policy initiatives, such as energy data-sharing frameworks, promote the responsible release of smart meter data to support research, innovation, and climate policy [26, 34]. While these frameworks advance the green energy transition, they

also increase the circulation of potentially sensitive data. Simultaneously, new regulatory instruments—such as the Digital Services Act [32] and the AI Act [35]—tighten constraints on data use and automated decision-making [59]. These intersecting dynamics create a tension between fostering innovation and ensuring privacy. Rising cybersecurity threats targeting the energy sector [27] further exacerbate the risks of unauthorized access and re-identification [73].

In light of this complex landscape, data controllers may be reluctant to share smart meter data. When they do, they often rely on aggregation-based PETs to achieve compliance [30]. One approach that offers a potential balance between privacy protection and analytical utility is microaggregation, which achieves k -anonymity by grouping similar consumption profiles and substituting individual consumption profiles with their group-level average [2]. While this reduces the likelihood of singling out and linkage, it does not eliminate the inference risk [70]. Distinct behavioral patterns may persist, allowing adversaries to **re-link pseudonymized records with anonymized groups** and complement them with publicly available external metadata.

Although some socio-demographic characteristics, such as sauna ownership, may already be inferred from raw pseudonymized load profiles, group-level linkage can enrich this knowledge, for example by enabling inferences about income levels or occupancy patterns with higher confidence, and by cross-validating behavioral signals across data releases. Such inferences, even at the group level, may be exploited for discriminatory profiling, targeted advertising, dynamic pricing by insurers, or other privacy-invasive practices, without requiring explicit household-level identification.

To better understand the effectiveness and limitations of microaggregation in this context, we systematically investigate the residual privacy risks it entails. We design a replicable methodology based on a publicly available smart meter dataset that has been anonymized using microaggregation. Our objectives are threefold: First, we assess the extent to which behavioral similarity enables cross-dataset linkability by simulating group membership inference attacks. Secondly, we evaluate the impact of microaggregation on data utility in the context of energy suppliers performing household-level, day-ahead load forecasting, by examining the effects of shifting from pseudonymised to anonymized datasets. Third, we jointly explore the trade-off between privacy protection and utility retention that emerges from applying microaggregation. Consequently, our study makes three key contributions:

- (1) We formulate and implement a group membership inference attack on micro-aggregated smart meter data. The attack simulates adversaries with varying degrees of expertise, employing unsupervised similarity-based matching and a two-stage majority voting mechanism to identify the most likely anonymized group of a household. This enables privacy risk quantification across different levels of k -anonymity.
- (2) We explore the impact of microaggregation on data utility by performing household-level day-ahead load forecasting using representative machine learning models. Our analysis presents empirical results on how increasing the anonymity parameter k reduces forecasting accuracy.

- (3) We present a replicable methodology, explained in Section 4, for evaluating utility and privacy that simultaneously assesses the success of inference attacks and forecasting performance. Our methodology identifies specific threshold values of k where privacy gains begin to outweigh utility losses. This insight serves as a guide for selecting appropriate levels of anonymization for responsible data sharing.

We organize the remainder of the paper as follows. Section 2 introduces key privacy concepts underlying this work. Section 3 reviews relevant literature. Section 4 outlines our research approach, and Section 5 defines the threat model guiding our analysis. Section 6 describes the experimental setup, while Section 7 presents and interprets the results from both utility and privacy perspectives. Finally, Sections 8 and 9 discuss broader implications, highlight future directions, and conclude the paper.

2 PRIVACY CONCEPTS

To support a clear understanding of the privacy risks and protection strategies discussed in this work, this section introduces foundational privacy concepts and anonymization techniques relevant to smart meter data.

Personally Identifiable Information (PII) includes not only direct identifiers such as names or identification numbers but also indirect or **quasi-identifiers**—attributes that may be linked to individuals through combinations or patterns in the data [52, 63]. In the context of electricity consumption, behavioral patterns embedded in load profiles can function as quasi-identifiers, e.g., enabling the re-identification of individuals even in the absence of explicit identifiers [18].

We understand **pseudonymization** in accordance with the General Data Protection Regulation (GDPR) regulation [33]. In the case of load profiles, pseudonymization involves the removal of direct identifiers (e.g., names, addresses) and the substitution of those with alphanumeric quasi-identifiers. Table 1 provides an illustrative example, where the attribute `Meter_id` serves as the quasi-identifier. This approach preserves the structure and resolution of the original data—particularly the values in the `kWh` column—while aiming to obscure direct links to identity.

Table 1. Pseudonymized load profiles of two individual households.

| Meter_id | Time | kWh |
|---------------|------------------|-------|
| Z12345 | 01.01.2025 00:30 | 0.346 |
| .. | .. | .. |
| Z12345 | 31.12.2025 23:30 | 0.004 |
| Z12346 | 01.01.2025 00:30 | 0.158 |
| .. | .. | .. |
| Z12346 | 31.12.2025 23:30 | 0.272 |

However, pseudonymization does not alter the underlying behavioral patterns, which may still allow for re-identification when combined with auxiliary data. Prior studies have shown that unique consumption behaviors can re-link records to individuals, even when direct identifiers are removed [13, 43].

In contrast, **anonymization** refers to a collection of techniques that irreversibly transform personal data to ensure individuals are no longer identifiable, either directly or indirectly [14, 33, 70]. Anonymization methods are typically classified as either **perturbative** or **non-perturbative**, depending on whether the original data values are altered [24].

Non-perturbative approaches, such as suppression and generalization, preserve the original values but reduce identifiability by removing or coarsening data granularity. For example, exact timestamps may be generalized into hourly intervals or suppressed entirely. These methods can be effective in applications where the temporal precision of the data is not critical. A prominent example is smart meter billing, where only the total energy consumption within a billing cycle (e.g., monthly or yearly billing cycle) is needed. In this context, precise consumption patterns or continuous load curves are not required, and non-perturbative anonymization may suffice to protect privacy without sacrificing utility.

Perturbative approaches, by contrast, introduce controlled modifications to the load profiles to obscure individual identities. These include techniques such as noise addition, data swapping, and synthetic data generation [25, 39]. Among these, **microaggregation** has emerged as particularly relevant for time-series datasets. Microaggregation operates by grouping k similar records—such as daily load profiles—and replacing individual values with the group’s average [22, 23]. This method retains the temporal structure of the data, which is essential for analytical tasks like day-ahead load forecasting, while still offering protection against re-identification. As expressed in Table 2, micro-aggregated data no longer represents individual behavior but captures collective patterns that can still be useful for operational tasks.

Table 2. Micro-aggregated load profiles of the households from Table 1.

| Group_id | Time | kWh |
|----------|------------------|-------|
| 1 | 01.01.2025 00:30 | 0.252 |
| .. | .. | .. |
| 1 | 31.12.2025 23:30 | 0.138 |

Importantly, microaggregation satisfies the principle of k -**anonymity**, meaning that each anonymized record is indistinguishable from at least $k-1$ others based on quasi-identifiers. This makes it a viable option in contexts where privacy protection and the retention of analytical utility—such as load forecasting—are required.

3 RELATED WORK

Following our load forecasting and data attacks scope, we present our explored related work. In Section 3.1, we motivate and briefly examine load forecasting as our utility case, highlighting its relevance, the sensitivity of the data it requires, and the surge of new forecasting methods leveraging smart meter data. However, we do not aim to provide a thorough review of the load forecasting literature. Instead, we use it as a representative application domain, drawing on well-established forecasting models already extensively explored in academia. In contrast, Section 3.2 focuses on relevant academic literature concerning privacy risks and attacks on load data and Section 3.3 clarifies our novelty.

3.1 Load Forecasting

Load forecasting plays a central role in smart grid operations, supporting tasks such as grid balancing, demand response, and planning for renewable integration. Recent advances in forecasting leverage both classical machine learning and deep learning to capture the temporal dynamics and behavioral variability in household-level consumption [55].

[50] propose a sequence-to-sequence recurrent neural network model for short-term load forecasting, illustrating the capability of deep learning to capture non-linear temporal dependencies in fine-grained consumption data. However, their approach operates on raw, non-anonymized load profiles without considering privacy-preserving data transformations, leaving open questions regarding forecasting performance under anonymization constraints. In contrast, our work explicitly assesses how forecasting accuracy is affected when using data anonymized via microaggregation, thereby addressing the impact of privacy-preserving interventions on model effectiveness.

[12] evaluate various machine learning models, including gradient boosting and support vector regression, for short-term residential load forecasting. They highlight the trade-offs between computational efficiency and predictive accuracy but, similar [50], conduct their analysis using fully accessible data without applying anonymization techniques. Our study extends this research by systematically quantifying the degradation in forecasting performance across different levels of microaggregation, providing insights for privacy-aware forecasting in smart grid applications.

[71] introduce an attention-based encoder-decoder architecture with bidirectional LSTM layers for multi-horizon short-term load forecasting, dynamically weighting historical and similar-day features to achieve state-of-the-art accuracy on public datasets. While their focus is on advancing forecasting performance through model sophistication, our work instead evaluates how strong standard architectures—both tree based and neural network-based—perform under privacy-preserving microaggregation. This approach allows us to analyze how anonymization impacts forecasting utility without further confounding factors from architectural advancements.

In summary, while prior work has demonstrated effective forecasting methodologies under conditions of complete data availability, our study explicitly investigates the forecasting performance of representative models on microaggregated smart meter data. This enables a quantifiable understanding of how privacy-preserving data sharing practices influence the utility of load forecasts within realistic smart grid scenarios.

3.2 Privacy Attacks on Load Data

Load profiles exhibit strong temporal and behavioral structures that make them highly distinctive on a per-household basis. Even in the absence of explicit identifiers, adversaries have exploited these patterns to successfully **re-link pseudonymized load profiles to individual households**. Early work [43] demonstrated that pseudonymized load profiles could be re-identified using behavioral anomalies and pattern matching, even under common mitigation strategies such as re-pseudonymization or reduced temporal resolution. [13] extended this approach by linking pseudonymized

fine-grained data with monthly aggregates, achieving full de-pseudonymization through iterative value matching. [57] introduced a method using feature-based similarity matching on weekly load profile snippets, applying majority voting to recover individual identities. Most recently, [16] employed deep learning-based embeddings and nearest-neighbor classifiers to re-identify households.

While prior research has primarily focused on re-identification attacks against pseudonymized load profiles, our study shifts the focus to a more privacy-preserving scenario involving explicitly anonymized smart meter data. Specifically, we investigate whether individual households can still be re-linked with anonymized groups through group membership inference. This extends the threat landscape from individual-level re-identification to a more complex variant of record linkage across datasets subjected to perturbative anonymization.

[21] examine re-identification risks under an anonymization model that relies on constrained permutation. In their approach, household identifiers are removed, but high-resolution load profiles remain intact. The adversary is assumed to have access to aggregate billing data, such as monthly or yearly consumption values, and attempts to re-establish the mapping between anonymized high-resolution records and known aggregates. By exploiting the inherent consistency between high-resolution data and aggregate totals, they demonstrate that high re-identification success is achievable, even without explicit identifiers. Importantly, their attack targets datasets where the temporal and behavioral fidelity of individual profiles is fully retained.

In turn, [8] investigate the privacy risks associated with **aggregation-based anonymization**, simulating scenarios where household load profiles are grouped and aggregated. Unlike [21], they focus on the risk of **group membership inference** rather than individual re-identification. Using a formal indistinguishability-based privacy framework, they quantify whether an attacker can determine if a known load profile is part of an aggregated group. Their experiments, based on Non-Intrusive Load Monitoring (NILM) datasets, reveal that certain high-consumption patterns (e.g., from electric vehicles or heating systems) remain detectable even when data is aggregated over 10–20 households. Notably, their analysis uses simulated load profiles derived from appliance-level traces, not real-world smart meter data.

3.3 Novelty

In contrast to the covered studies in Section 3.2, we evaluate privacy risks under microaggregation, a perturbative anonymization technique that replaces individual profiles with the average of k similar profiles. In our implementation, groups of similar records are identified based on their distance to the group average to ensure low within-group variance while achieving k -anonymity. Our analysis builds upon and extends [8] membership inference model in two important ways. First, our attacker model operates over a **many-to-many matching problem**: rather than guessing whether one specific load profile is part of one aggregated group, we attempt to re-link multiple pseudonymized household profiles to their corresponding aggregated groups based on behavioral similarity. Second, we use real-world household-level smart meter data (from the Low

Carbon London dataset), thereby capturing a richer variety of usage patterns and privacy risks.

From a terminological standpoint, our attack constitutes a record linkage mechanism designed to achieve **group membership inference**. While traditional membership inference attacks (e.g., in machine learning) are formulated as binary classification problems [60], our setting involves distance-based matching across two datasets: one containing pseudonymized individual profiles, the other containing anonymized group averages. Although the method is grounded in similarity-based record linkage, the goal remains group-level association – identifying to which anonymized group each individual most likely belongs. This formulation is more complex than classical binary group membership decisions, as it operates without direct identifiers and across multiple households and groups simultaneously.

Finally, our work contributes a utility-privacy perspective, absent from prior studies. While [21] and [8] focus exclusively on privacy leakage, we assess how increasing the microaggregation parameter k simultaneously affects both privacy (via group membership inference success) and utility (via household-level load forecasting accuracy). This analysis, combining perturbative anonymization, unsupervised similarity-based linkage, and forecasting-based utility assessment, offers new insights into the residual privacy risks and trade-offs involved in sharing micro-aggregated smart meter data.

4 RESEARCH APPROACH

To support informed decisions about privacy-preserving data sharing, we develop a structured research approach that assesses both the risks and utility of micro-aggregated smart meter data. This section outlines the underlying problem and research questions, followed by a replicable methodology that guides our analysis.

4.1 Problem Statement and Research Objectives

This work investigates whether load profiles, once anonymized through microaggregation [23], still permit the inference of individual household membership within aggregated groups. While microaggregation is designed to achieve k -anonymity by replacing individual load profiles with the group average of k similar households, it remains unclear whether distinctive behavioral patterns persist in the aggregated data. If such patterns survive anonymization, adversaries may be able to re-link pseudonymized household records with their corresponding micro-aggregated groups, posing privacy risks even when individual identifiers are removed.

The central objective of this study is to assess how well micro-aggregated load profiles balance analytical utility with privacy protection. We examine the extent to which weekly consumption patterns enable group membership inference, and how increasing the aggregation parameter k influences both the success of such inference attacks and the utility of the data for downstream tasks such as household-level load forecasting. To structure this investigation, we pose the following research questions:

- (Q1) To what extent can attackers infer the anonymized group membership of individual households from micro-aggregated load profiles, and how does this privacy risk

evolve with increasing values of the aggregation parameter k ?

(Q2) How does increasing the anonymization level k impact the utility of anonymized load profiles for household-level day-ahead forecasting?

(Q3) At which values of the microaggregation parameter k does a measurable decline in group membership inference success (privacy gain) coincide with an acceptable loss in forecasting accuracy (utility degradation), and how can this trade-off be used to inform parameter selection for privacy-preserving smart meter data sharing?

4.2 Methodology

Figure 1 illustrates the used replicable methodology, which comprises five steps designed to systematically evaluate privacy risks and utility in anonymized load profiles.

1) Data Selection and Preparation: This study is based on the publicly available Low Carbon London (LCL) dataset, which serves as a widely recognized benchmark in smart meter analytics and privacy research [65]. The dataset is pseudonymized, meaning that household identifiers have been replaced with alphanumeric codes while preserving load profiles recorded at 30-minute intervals. While our evaluation is conducted using the LCL dataset, the methodology is designed to be generalizable to other time series datasets that contain regular interval load measurements (e.g., 15- or 30-minute resolution).

2) Anonymization (Microaggregation): To simulate a realistic privacy-preserving data release by an energy provider, we apply anonymization to the pseudonymized smart meter data using microaggregation. This process groups similar household load profiles and replaces individual consumption traces with group-level averages, thereby obscuring household-specific patterns. The resulting anonymized load profiles serve as the basis for evaluating both privacy risk and data utility in downstream applications.

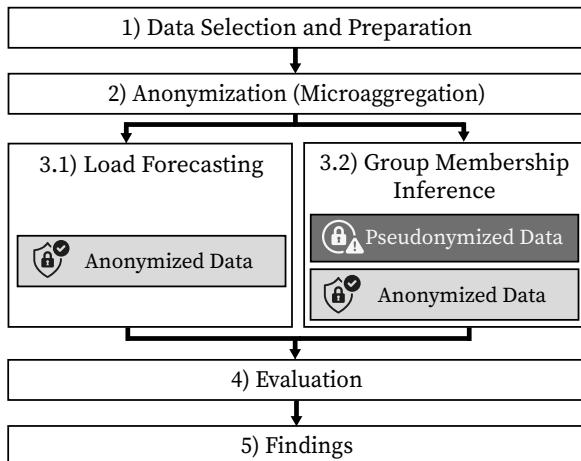


Fig. 1. Overview of the methodological pipeline. The process spans data selection, anonymization, scenario execution, evaluation, and analysis.

3) Scenario Implementation: We implement two analytical scenarios to assess the utility and privacy implications of microaggregated load profiles. First, in the **household-level load forecasting** scenario, we train predictive models on anonymized group profiles and evaluate predictive performance on individual household consumption (see Section 6.3). Second, in the **group membership inference** scenario, we simulate a privacy attack using distance-based record linkage to link weekly pseudonymized household load profiles with the most similar micro-aggregated group profiles (see Section 6.4).

4) Evaluation: We evaluate both scenarios independently to assess their respective outcomes. Forecasting performance is measured using household-level prediction error (Section 7.1), while privacy risk is quantified by the success rate of the group membership inference attack (Section 7.2). In a second step, we jointly analyze both results to examine the trade-off between utility and privacy across varying levels of anonymization (Section 7.3).

5) Findings: By systematically varying the anonymization parameter k , we explore how the relationship between forecasting utility and privacy risk evolves. While the resulting trade-off curve combines two distinct metrics, prediction error and group membership inference success, it serves as an initial attempt to visualize their interaction. Rather than prescribing specific values for k , the analysis offers a qualitative indication of regions where privacy protection improves without sharply compromising utility, providing a basis for future, more rigorous assessments.

5 THREAT MODEL

To systematically assess privacy risks in our setting, we adopt a formal threat modeling approach [7, 67]. This allows us to clearly define the adversary's objective, the underlying assumptions, and the knowledge and capabilities required to launch a group membership inference attack.

5.1 Adversary's Knowledge and Capabilities

The adversary operates under a closed-world assumption, meaning the pseudonymized and anonymized datasets refer to the same set of households without any extraneous households. Both datasets cover an identical time frame.

We assume that the attacker is aware of the anonymization method, microaggregation, and either knows or can reasonably estimate the group size parameter k , which typically ranges from 5 to 200. This assumption is realistic in regulated domains like energy systems, where privacy-preserving mechanisms may be disclosed as part of legal compliance or standardization [72].

Figure 2 illustrates the group membership inference scenario under consideration. On the left, individual pseudonymized household load profiles are collected and processed by a data aggregator, such as an energy supplier. These profiles undergo anonymization via microaggregation, before being stored or potentially shared.

The resulting anonymized dataset is either shared—intentionally for research or commercial purposes—or leaked through data breaches or passive interception (eavesdropping). In both cases, the adversary is assumed to gain access to anonymized load profiles,

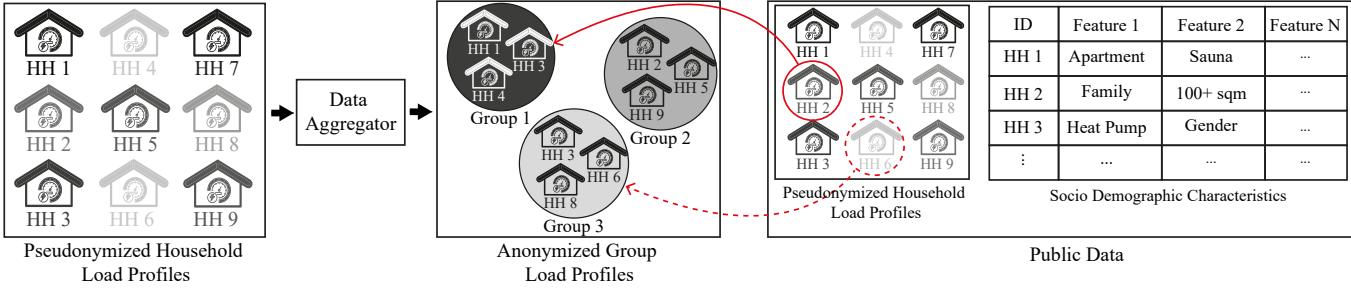


Fig. 2. Illustration of the adversary scenario for group membership inference. Pseudonymized household load profiles are collected and anonymized via microaggregation. The adversary leverages publicly available pseudonymized data, enriched with socio-demographic information, to re-link individual household profiles to anonymized groups. The dashed arrow indicates a successful inference, where the adversary correctly assigns Household 6 to Group 3, thereby revealing its group membership.

which serve as the target of the attack. The adversary's goal is to re-link individual pseudonymized profiles to their corresponding aggregated group, thereby performing a group membership inference attack. This process is depicted by the red arrows in Figure 2. The dashed arrow indicates a successful inference, where the targeted household is correctly linked to its group, while the solid arrow represents an attempted inference where the household is not part of the examined group.

To perform this attack, the adversary draws upon two key data sources: (i) Pseudonymized load profiles with auxiliary information: These publicly available datasets include household-level load profiles and auxiliary socio-demographic attributes. Although direct identifiers are removed through pseudonymization, associated metadata such as income group, household size, dwelling type, or appliance ownership often remains intact. Such data is commonly released in open datasets [44, 69], making this a realistic assumption.

(ii) Anonymized load profiles obtained through microaggregation: The adversary is assumed to have access to a micro-aggregated dataset—comprising group-level average load profiles—for a duration of **at least one week**. This time frame reflects a realistic attack window, as short-term data exposures may occur through insider leaks or passive eavesdropping on communication infrastructures [73]. Prior work [5, 41, 57] have demonstrated that even weekly load profiles can retain sufficient behavioral information to compromise individual privacy, underscoring the feasibility of inference attacks based on limited temporal time frames.

To simulate varying levels of adversarial capability, we implement three attacker models: (i) a baseline attacker that employs simple statistical features such as means and variances; (ii) a domain-informed attacker that extracts behavioral descriptors based on established features from prior work [4, 5]; and (iii) a representation learning-based attacker that utilizes autoencoders, whose architectures are designed to reflect domain knowledge of temporal and behavioral load patterns. All attacker models operate in a fully unsupervised setting and **do not access ground-truth mappings during inference**. Ground-truth labels are solely used in the evaluation phase to assess attack performance, as described in Section 6.

5.2 Adversary's Objective

The adversary aims to re-link publicly available pseudonymized load profiles to their corresponding groups in a dataset anonymized through microaggregation. The adversary's objective is to determine the group g_h to which the target household h belongs.

The adversary has access to two key data sources: a pseudonymized dataset $D = \{c_1, c_2, \dots, c_N\}$, where c_h is the pseudonymized load profile of household $h \in \{1, \dots, N\}$, and the micro-aggregated dataset $\tilde{D} = \{\tilde{c}_1, \tilde{c}_2, \dots, \tilde{c}_{N/k}\}$ ($k \mid N$), where each \tilde{c}_g represents the average load profile of a group of k households. Microaggregation groups D into N/k groups, where $g \in \{1, \dots, N/k\}$. Each household h belongs to exactly one group g_h .

We formalize the adversary's objective through a security game referred to as the **group membership inference game** [8]. The game is played between two parties named challenger Ch and adversary Adv . The challenger is an abstract entity that represents all parties who are concerned about their privacy. The adversary is a party who acts on behalf of entities who aim at violating privacy.

- (1) The adversary Adv is given the datasets D and \tilde{D} .
- (2) A random household h is chosen by Ch and given to Adv .
- (3) Using the obtained data, the adversary Adv tries to determine g_h . His guess is denoted as g' .
- (4) The adversary outputs g' and wins if and only if his guess is correct, i.e., $g' = g_h$.

Intuitively, privacy is broken if the adversary has a higher winning probability than random guessing, i.e.,

$$\Pr[g' = g_h] > \frac{1}{N/k} \quad (1)$$

Note that the game does not specify how the adversary analyzes the given data as these are details of the attacks that depend on the attacker types described before. These specific parts of the attack are described in Section 6.

5.3 Threat Scope and Impact

Although the proposed attack does not reveal individual household identities, it compromises privacy by linking pseudonymized load profiles to their corresponding micro-aggregated groups. This linkage enables adversaries to cross-reference behavioral load data with auxiliary socio-demographic information, thereby transforming anonymized data into enriched sources of sensitive inference.

This enrichment data could be used to amplify vulnerabilities inherent in group-based anonymization schemes [3]: The fact that microaggregation selects similar households for the clustering of groups and that microaggregation retains correlations between quasi-identifiers and sensitive attributes could be amplified by this enrichment data.

When an adversary possesses external knowledge that links specific consumption patterns to known socio-demographic traits—such as regular overnight charging associated with electric vehicle (EV) ownership or high evening usage linked to affluent households—they can use this information to draw inferences about group members. Successful linkage of pseudonymized profiles to micro-aggregated groups allows these behavioral patterns to serve as proxies for inferring sensitive characteristics, even in the absence of direct individual identification.

For example, associating a pseudonymized household with a group showing elevated late-night energy consumption and minimal daytime use could suggest shift work patterns or illicit indoor agriculture. Likewise, groups characterized by consistently high summer usage may indicate energy-intensive appliances like air conditioning during specific times of the day or pool pumps, which can be proxies for household affluence. Such inferences, even at the group level, may be exploited for discriminatory profiling, targeted advertising, dynamic pricing by insurers, or other privacy-invasive practices, without requiring explicit household-level identification.

This attack model challenges the sufficiency of microaggregation as a standalone anonymization technique and suggests that behavioral fingerprints may survive aggregation, particularly when combined with publicly available auxiliary data and machine learning techniques.

6 EXPERIMENTAL SETUP

This section outlines the experimental setup used to evaluate the utility and privacy implications of micro-aggregated load profiles. We begin by describing the dataset and preprocessing steps, followed by the microaggregation method applied for anonymization. Subsequently, we detail the forecasting models used to assess data utility and present the attack methodology employed for group membership inference.

6.1 Data Description

The dataset used for our experiments is based on the Low Carbon London dataset, provided by UK Power Networks, from which we consider 4342 household load profiles within the London area from November 2011 to February 2014 [17]. To reduce computational demands for the experiments, we randomly select a subset of 1,000 households from January 1, 2013, to December 31, 2013.

The data provides half-hourly electricity consumption readings, resulting in 48 values per day per household. To contextualize the selected sample, we compute the annual energy consumption per household and find an average of 3688 kWh and a median of 3013 kWh, aligning well with typical residential usage patterns in urban UK settings during that period [66].

6.2 Data Protection via Microaggregation

As already introduced in Section 2, we apply microaggregation, which partitions load profiles into clusters of at least k similar households and replaces each individual profile with the group average. This structure-preserving transformation reduces re-identification risk by ensuring that no household is distinguishable from fewer than $k - 1$ others [23].

Unlike naive grouping strategies, microaggregation relies on **similarity-based clustering** to preserve utility-relevant properties in the data. In this study, we employ the Maximum Distance to Average Vector (MDAV) algorithm to construct anonymized clusters. MDAV iteratively selects the record farthest from the global average and groups it with its closest neighbors, thereby **minimizing intra-cluster variance**. This approach is particularly effective for preserving the statistical and temporal characteristics of load profiles, which are crucial for downstream tasks such as forecasting [2].

While advanced variants such as DFTMicroagg [2] apply frequency-domain transformations, most notably the Discrete Fourier Transform (DFT), to enhance anonymity, we choose standard microaggregation based on the MDAV algorithm. This choice is motivated by MDAV's ability to balance computational efficiency with the preservation of temporal structures in the original data, an essential property for maintaining forecasting performance. Unlike DFT-based approaches, which may introduce spectral artifacts or distortions, MDAV operates directly in the time domain, ensuring the integrity of sequential consumption patterns. Moreover, MDAV is widely regarded as the most commonly used microaggregation algorithm in practice, particularly in tools like the sdcMicro package [64].

To systematically explore the utility-privacy trade-off, we construct a series of anonymized datasets by applying microaggregation at varying levels of the privacy parameter k . Formally, we define a finite set of privacy levels $k = \{5, 10, 25, 50, 100, 200\}$, where lower values of k (e.g., $k = 5$) correspond to minimal privacy protection and higher values (e.g., $k = 200$) represent stronger anonymity guarantees. These values allow us to observe the impact of aggregation strength on both utility and privacy.

6.3 Load Forecasting for Utility Assessment

To assess the utility of micro-aggregated load profiles, we focus on household-level day-ahead load forecasting. This task serves as a representative benchmark for assessing how well anonymized data supports typical energy analytics. To ensure temporal robustness and minimize bias from specific time windows, we implement a rolling cross-validation strategy using 31-day training windows, each shifted forward by one day.

We select five forecasting models including two deep learning models (NBEATS [54] and NHITS [9]) capable of capturing non-linear temporal structures; two tree-based ensemble methods (XG-Boost [11] and LightGBM [45]) known for their accuracy and efficiency; and a linear regression model serving as an interpretable baseline. This diverse model selection allows us to assess how different algorithmic approaches respond to increasing levels of data anonymization. Each model is trained on micro-aggregated profiles corresponding to varying k -anonymity levels. The models are then asked to produce a forecast based on the original, non-anonymized household-level load data. The models are then evaluated based on this forecast. This can be understood as transfer learning from models trained on anonymized data to non-anonymized data. To reduce computational overhead across multiple anonymization configurations, we use default model parameters from standard libraries such as scikit-learn, foregoing hyperparameter optimization.

Forecasting accuracy is measured using multiple error metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Squared Error (MAE), and Symmetric Mean Absolute Percentage Error (SMAPE) [61]. We include SMAPE in place of MAPE to address the known issue of instability when actual consumption values approach zero—a common occurrence in household load profiles. All models produce 48 half-hourly predictions per day per household, for both the raw and anonymized datasets. A detailed discussion of forecasting results across varying values of k is presented in Section 7.1.

6.4 Attacking Micro-aggregated Load Profiles

The objective of the group membership inference attack is to assess whether individual household load profiles can be correctly linked to their corresponding anonymized group load profiles using distance-based record linkage. To achieve this, we develop a similarity-based matching methodology that is trained solely on micro-aggregated load data and subsequently applied to pseudonymized household profiles for inference.

The attack methodology consists of a four-stage pipeline: (i) data preparation, (ii) feature extraction, (iii) similarity matching, and (iv) evaluation. The first three stages are illustrated in Figure 3. Notably, the entire process operates in a fully **unsupervised** setting—no ground-truth labels are accessible during stages (i) to (iii). Ground-truth information is used exclusively in stage (iv) to evaluate the success of the group membership inference attack.

6.4.1 Data Preparation. To enable a consistent and meaningful comparison between pseudonymized household profiles and micro-aggregated group profiles, both datasets are aligned to share the same temporal and structural format. The data spans 51 full weeks, covering the period from January 1 to December 31, 2013. The final calendar week (week 52) has been excluded, as it lacks a complete 7-day record. Each load profile, whether individual or aggregated, is segmented into **weekly intervals**, following the methodology of [57].

Consequently, each pseudonymized household profile h and each micro-aggregated group profile g is segmented into 51 weekly snippets. Such a snippet of household h and week w is denoted as $c_{h,w}$ and consists of $7 \cdot 24 \cdot 2 = 336$ time-aligned measurements. The

same procedure is done for each of the N/k group profiles whose weekly snippets are denoted as $c_{g,w}$. While the pseudonymized load profiles retain **individual** half-hourly consumption values, the micro-aggregated load profiles contain the corresponding **averaged** measurements across groups of k households. The resulting data structures are illustrated in the first step (1. Data Preparation) of the attack methodology, depicted in Figure 3.

6.4.2 Feature Extraction. To simulate varying levels of adversarial capability as described in Section 5.1, we implement three distinct feature extraction strategies: (i) basic statistical descriptors—such as mean, standard deviation, minimum and maximum—computed for each weekly load profile; (ii) domain-specific features derived from 35 numerical descriptors proposed by [4, 5]; and (iii) automated feature learning using deep neural networks, specifically a Recurrent Autoencoder (RAE) and a Convolutional Autoencoder (CAE). The resulting features are denoted as $x_{h,w}$ for pseudonymized household and $x_{g,w}$ for the corresponding group load profiles.

Our primary representation is based on the learned embeddings generated by the RAE and CAE models. The design of these models was guided by the behavioral structure captured in the [4, 5] feature set. Specifically, kernel sizes in the CAE and the LSTM architecture in the RAE were tailored to reflect key aspects such as temporal dynamics and intra-day consumption ratios. This architecture design aims to simulate characteristic load behaviors directly within the learned representation. The final embedding dimensionalities are $p = 48$ for the CAE and $p = 32$ for the RAE. Detailed architectural configurations for the CAE and RAE are provided in Table 5 and Table 6, respectively.

While domain-specific features provide high interpretability and remain useful for analysis, the automated feature learning approaches offer superior performance in capturing nuanced consumption patterns that may persist even after microaggregation.

6.4.3 Similarity Matching (k nn). To infer group membership from pseudonymized household load profiles, we implement a **similarity-based k -nearest neighbor** approach using the extracted p -dimensional feature vectors. For each weekly pseudonymized household profile $x_{h,w}$, we compute the Euclidean distance to all anonymized weekly group profiles $\tilde{x}_{g,w} \in \tilde{D}$, where $g = 1, \dots, N/k$ and $w = 1, \dots, 51$. The n_{nn} closest group-week profiles, i.e., those with the smallest pairwise distances, are selected as candidate matches.

The matching process involves a **two-stage majority voting procedure** that is formally described in Algorithm 1. In the first stage, each household-week snippet $x_{h,w}$ is assigned to the group that appears most frequently among its n_{nn} nearest neighbors. In case of ties, we select the group with the smallest average distance to $x_{h,w}$. This detail is omitted in Algorithm 1 for sake of simplicity.

In the second stage, a final group prediction g'_h is determined for each household by majority voting across the weekly group assignments $g'_{h,1}, \dots, g'_{h,51}$. The same distance-based tie-breaking strategy is applied if necessary.

Table 3 presents an example of top-3 nearest neighbors ($n_{nn} = 3$) and resulting group assignments for selected snippets for a single household h . For a better understanding of the entire two-stage majority voting process, Figure 3 includes a simplified feature space

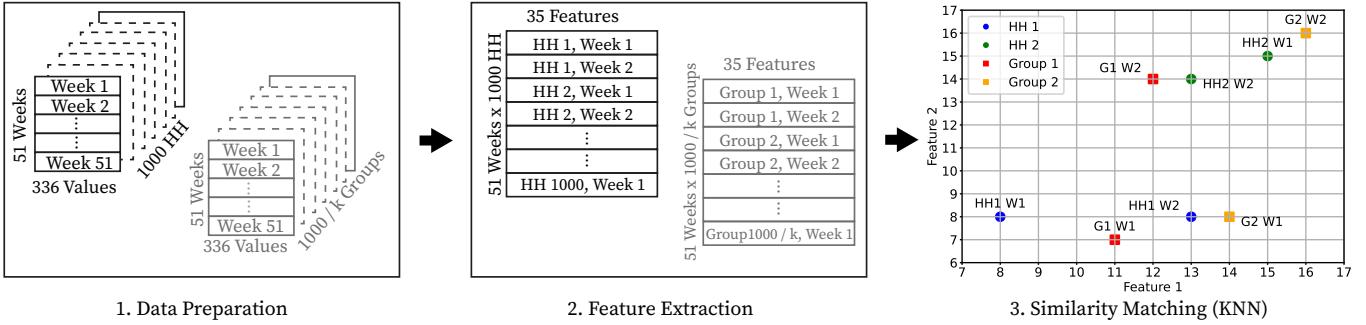


Fig. 3. Overview of the re-linking methodology comprising three stages: 1. Data preparation, 2. Feature extraction, and 3. Similarity matching via k-nearest neighbors. The figure illustrates the corresponding data structures at each stage. The feature extraction step highlights the interpretable, domain-specific case. A simplified two-dimensional example in stage 3 visualizes two weekly snippets from two households and two groups.

Algorithm 1 Similarity Matching Algorithm

Require: Feature representations $\{x_{h,w}\}_{h=1,\dots,N; w=1,\dots,51}$ of pseudonymized load profiles, feature representations $\{\tilde{x}_{g,w}\}_{g=1,\dots,N/k; w=1,\dots,51}$ of aggregation groups, privacy parameter k

- 1: **for** each household h in $\{1, \dots, N\}$ **do**
- 2: **for** each week w_1 in $\{1, \dots, 51\}$ **do**
- 3: Set $x = x_{h,w_1}$
- 4: Find the n_n nearest neighbors of x from all $51 \cdot N/k$ feature representations $\{\tilde{x}_{g,w}\}$.
- 5: Denote the groups g from this set as $\mathcal{N}(x)$
- 6: Voting 1: Set $g'_{w_1} = \arg \max_{g \in \{1, \dots, N/k\}} \sum_{g_i \in \mathcal{N}(x)} \mathbb{1}(g_i = g)$
- 7: /* Handling of ties omitted for simplicity */
- 8: **end for**
- 9: Voting 2: Set $g'_h = \arg \max_{g \in \{1, \dots, N/k\}} \sum_{w_1 \in \{1, \dots, 51\}} \mathbb{1}(g'_{w_1} = g)$
- 10: **end for**
- 11: **Output:** Assigned aggregation groups g'_1, \dots, g'_N for all N households

(right panel), illustrating two weeks from two households and their corresponding group profiles.

Table 3. Example illustrating the similarity matching algorithm using $n_n = 3$. It shows the top-3 nearest micro-aggregated weekly snippets (neighbors) for each week of a household h . The 51 groups g'_w are found by majority voting among the first (group) indices along the rows. The final group choice g'_h for household h is found by majority voting among these weekly group winners g'_w along the last column.

| Snippet | Neighbor 1 | Neighbor 2 | Neighbor 3 | g'_w |
|------------|--------------------|---------------------|--------------------|----------|
| $x_{h,1}$ | $\tilde{x}_{5,7}$ | $\tilde{x}_{5,3}$ | $\tilde{x}_{22,2}$ | 5 |
| $x_{h,2}$ | $\tilde{x}_{5,9}$ | $\tilde{x}_{67,50}$ | $\tilde{x}_{5,48}$ | 5 |
| \vdots | \vdots | \vdots | \vdots | \vdots |
| $x_{h,51}$ | $\tilde{x}_{42,6}$ | $\tilde{x}_{42,2}$ | $\tilde{x}_{8,44}$ | 42 |
| g'_h | | | | 5 |

6.4.4 Privacy Risk Evaluation. To evaluate group membership inference performance at the household level, we leverage the ground-truth mapping between each household h and its corresponding micro-aggregated group g_h . This mapping is only used during this evaluation phase. For each anonymity level k , we compare the final predicted group assignment g'_h obtained from the two-stage majority voting procedure, with the true group label g_h . This allows us to quantify how effectively the attack infers group membership under varying levels of anonymization.

For each k -anonymity level, we compute the **Attack Success Rate (ASR)** as the empirical estimator of the winning probability of the adversary game shown in Equation 1:

$$\text{ASR} = \frac{1}{N} \sum_{h=1}^N \mathbb{1}(g'_h = g_h) \quad (2)$$

Here, N denotes the number of households and $\mathbb{1}(\cdot)$ is the indicator function that returns 1 when the predicted group matches the true group, and 0 otherwise. In practice, this corresponds to computing `accuracy_score` from the `scikit-learn` library for each k .

However, the anonymization parameter k directly affects the number of possible candidate groups $|G|$, which in turn influences the difficulty of the classification task: as k increases, the number of candidate groups decreases, and the likelihood of correct assignment rises purely by chance. This makes raw accuracy an overly optimistic indicator of privacy risk at higher k values. To account for this effect, we normalize the observed accuracy by the baseline accuracy expected from random guessing. This normalized metric, often referred to as the **Relative Attack Success Rate (RASR)**, is widely used in literature to evaluate group membership inference attacks [42]. In our setting, we adapt this metric to reflect a multi-instance matching scenario, where multiple pseudonymized household profiles are re-linked to multiple anonymized groups.

$$\text{RASR} = \frac{\text{ASR}}{1/|G|} \quad (3)$$

An RASR of 1 corresponds to random guessing, whereas values greater than 1 indicate an elevated privacy risk, reflecting a higher-than-random success in group membership inference.

7 RESULTS

In this Section the impact of microaggregation on utility and privacy across varying levels of k -anonymity is evaluated: Section 7.1 reports forecasting accuracy for different models trained on aggregated data, while Section 7.2 analyzes the effectiveness of group membership inference attacks. Finally, both perspectives are combined in Section 7.3 to explore the emerging trade-off between data utility and privacy protection.

7.1 Utility: Load Forecasting

We evaluate data utility by assessing how microaggregation affects the accuracy of household-level day-ahead load forecasts. In order to explore possible dependencies on the model and the evaluation criterion, five forecasting models and five standard error metrics are used as described in Section 6.3. Since the choice of error metric illustrates only a minor influence on the results, we present the MAE as a representative metric in Figure 4. A complete overview of all metrics is provided in Appendix A, illustrating some statistics in Table 4.

Among the evaluated models, LGBM consistently delivers the strongest performance across most anonymization settings starting at 0.109 kWh for $k = 5$, followed closely by XGBoost that starts at slightly higher at 0.115 kWh for $k = 5$ the development as k increases is not as sharp as for LGBM. In contrast, the neural architectures NHITS and NBEATS show greater sensitivity to anonymization, with steeper performance degradation as the k parameter increases. This suggests that tree-based models are more robust to the distortions introduced by microaggregation.

As expected, forecasting performance declines across all models as the anonymity level k increases. This degradation is particularly evident when comparing performance relative to a lower-privacy baseline ($k = 5$). In high-privacy settings (e.g., $k = 200$), models such as Linear Regression and XGBoost exhibit over 10% performance loss in several error metrics.

7.2 Privacy: Membership Inference

We assess privacy by evaluating the success rate of group membership inference attacks, measuring how well adversaries can re-link pseudonymized household profiles to their corresponding micro-aggregated groups. As described in Section 5, we simulate attackers with varying levels of domain knowledge, using four distinct feature extraction strategies. Figure 5 illustrates the performance of these strategies across different anonymization levels. The left panel displays the group membership classification accuracy, while the right panel depicts the corresponding RASR, which normalizes accuracy by the success probability of random guessing, as described in Section 6.4.4.

For classification accuracy, we observe an overall increasing trend across higher k values due to the decreasing number of candidate groups, which naturally raises the likelihood of correct assignment by chance. However, a notable exception occurs in the transition from $k = 5$ to $k = 10$, where accuracy drops for all feature representations except the handcrafted approach.

The RASR metric, which corrects for this chance-level bias, reveals a consistent decline as k increases, indicating that stronger

anonymization reduces the effectiveness of group membership inference. Among the models, the CAE achieves the highest RASR across all k values, particularly in the low- k range ($k < 25$), where it significantly outperforms all other representations. The RAE and the domain-specific features provide moderate performance, while the handcrafted features remain close to random guessing.

To ensure robustness of the results, we test multiple neighborhood sizes $n_{nn} \in \{3, 5, 7, 9, 11\}$ for the similarity-based matching procedure and observe only minor variation in performance. For lower values of k (e.g., $k = 5$ and $k = 10$), smaller neighborhood sizes such as $n_{nn} = 3$ and 5 perform slightly better. We also evaluate cosine and Manhattan distances as alternatives to Euclidean distance, but find their results to be largely comparable, and therefore rely on Euclidean distance throughout the analysis.

7.3 Utility-Privacy Trade-off

We assess the trade-off between utility and privacy across varying levels of anonymization by jointly analyzing the performance of the most effective forecasting and attack models. As established in Section 7.1 and Section 7.2, LGBM demonstrates the highest forecasting accuracy, while the CAE yields the strongest group membership inference results. Figure 6 visualizes the trade-off between these two objectives. Forecasting utility is quantified via MAE on the left y-axis (blue), while privacy risk is represented by the RASR on the right y-axis (green). The x-axis spans the tested $k \in \{5, 10, 50, 100, 200\}$, illustrating how increasing anonymization impacts both objectives.

While the MAE gradually increases with growing k , the absolute values remain low across the full range, from approximately 0.110 to 0.124 kWh, indicating only a moderate degradation in predictive accuracy. In contrast, the RASR reveals a substantially steeper decline with increasing k , dropping from over 80 at $k = 5$ to below 30 at $k = 10$, and continuing to fall thereafter. This contrast in relative change highlights that substantial privacy gains can be achieved even at small aggregation levels, while utility losses remain limited. A complementary plot showing the relative change in MAE compared to the baseline ($k = 5$) is provided in Appendix C, Figure 9, offering an alternative perspective on the trade-off.

Overall, the plot illustrates that a substantial improvement in privacy protection—particularly in the transition from $k = 5$ to $k = 10$, can be achieved with only a minor reduction in forecasting performance, highlighting this region as a potential sweet spot in the utility-privacy trade-off.

8 DISCUSSION

Given our experimental setup, our findings illustrate that microaggregation can substantially mitigate the risk of group membership inference attacks on smart meter data, particularly as the anonymity parameter k increases. This protection becomes especially noticeable between $k = 5$ and $k = 10$, where the attack success rate drops sharply across all attacker types. Importantly, this privacy improvement occurs while forecasting utility remains relatively stable, particularly for robust models such as LGBM. This observation highlights a potential operational range in which a reasonable trade-off between utility and privacy can be achieved without compromising either objective substantially.

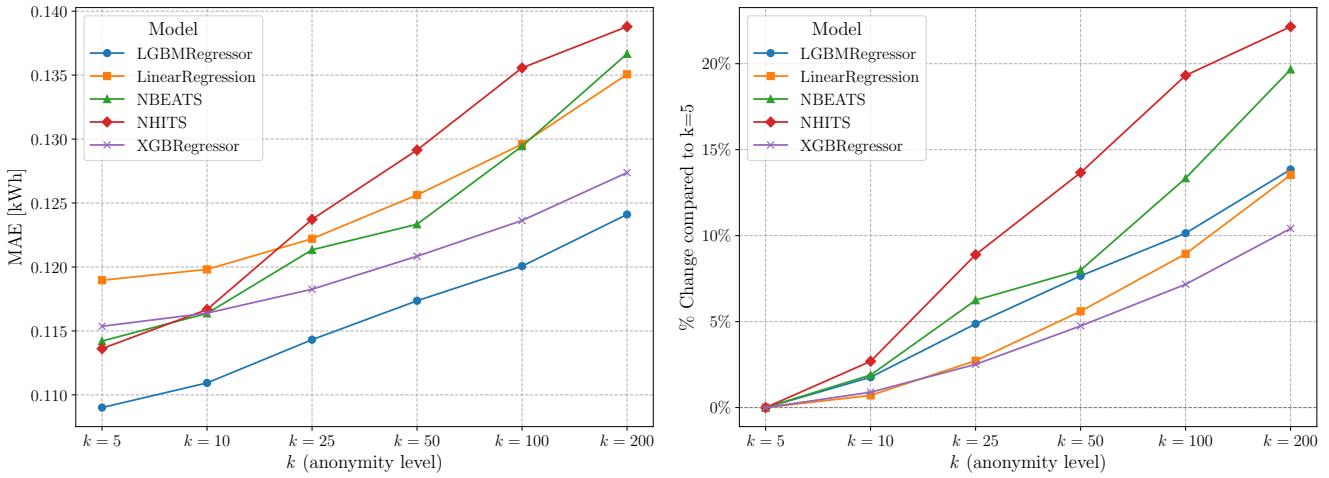


Fig. 4. Household-level day-ahead forecasting performance across varying anonymity levels. The left panel depicts the MAE for each model across increasing k . The right panel displays the relative performance degradation (% change in MAE) compared to the baseline scenario of $k = 5$.

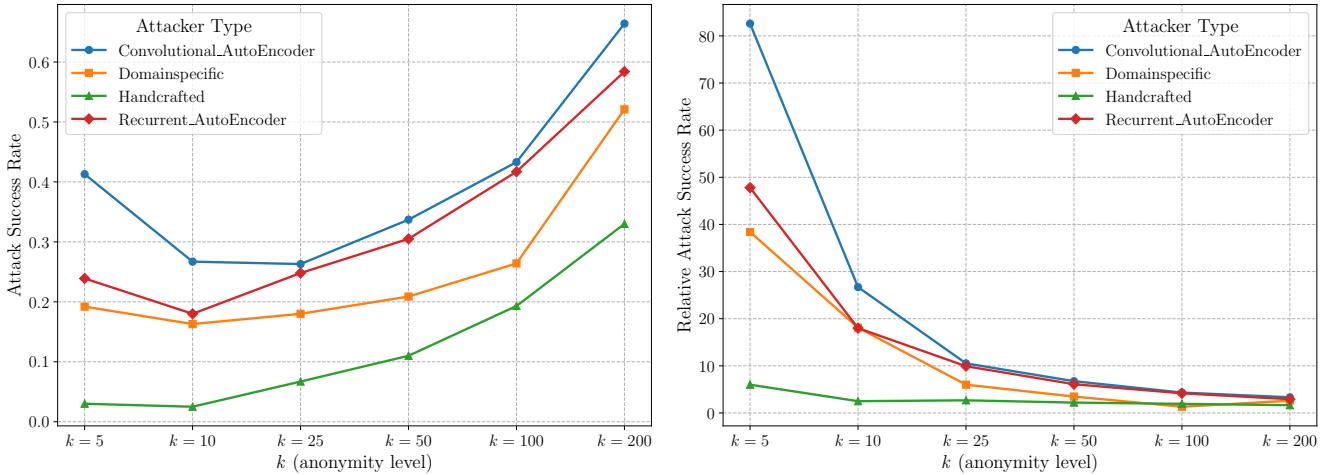


Fig. 5. Membership inference performance across varying k -anonymity levels. The left panel depicts the ASR for each adversary model across increasing k . The right panel displays the RASR.

However, the analysis also reveals that low aggregation levels (e.g., $k = 5$ and $k = 10$) offer limited protection against more sophisticated attacks. In this regime, deep learning models, especially the CAE, maintain high inference performance, although these values are relatively low for values between $k = 10$ and $k = 25$. The CAE's architecture, inspired by domain-specific features, appears to capture persistent behavioral patterns that remain even after aggregation. This highlights the inherent privacy risk that arises when structural patterns in load data are preserved for the sake of analytical utility.

One central reason for this persistence lies in the design of the microaggregation algorithm itself. MDAV, as a structure-preserving

method, minimizes intra-cluster variance by clustering behaviorally similar load profiles. While this benefits downstream applications such as forecasting, it also facilitates inference attacks by retaining patterns critical for distinguishing individual households, because the fewer groups exist, the easier it is, similar to the opposite case, as groups still retain patterns.

In Figure 7, a household with distinctive temporal patterns, most notably consistent low consumption between 4:00 and 6:00 a.m., and three unique dips occurring at the end of March, mid-July, and mid-August, is consistently and correctly re-linked to its anonymized group over all k -values. These unique temporal patterns, likely

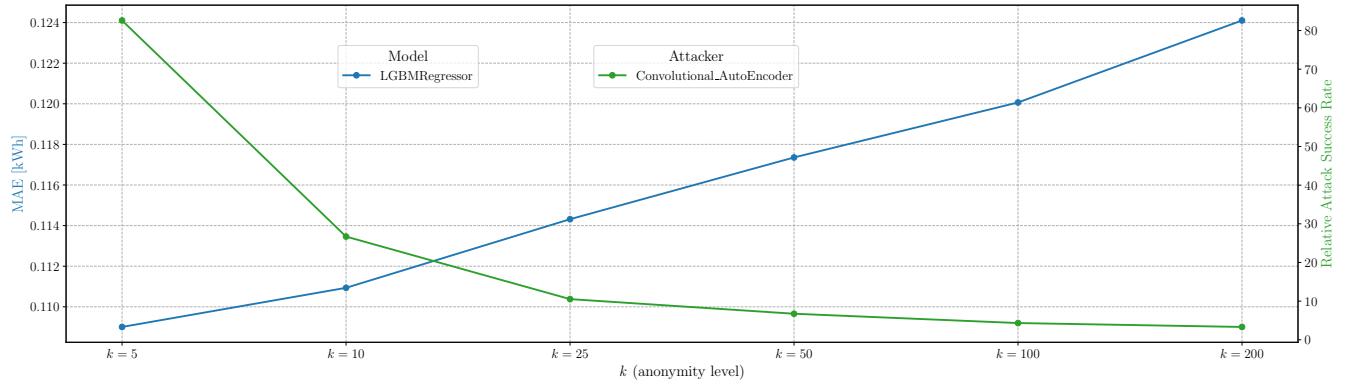


Fig. 6. Trade-off between utility and privacy for different levels of k -anonymity. Utility is quantified by the forecasting error (MAE, left y-axis, shown in blue) using a LGBM Regressor, while privacy is measured by the Relative Attack Success Rate (RASR, right y-axis, shown in green).

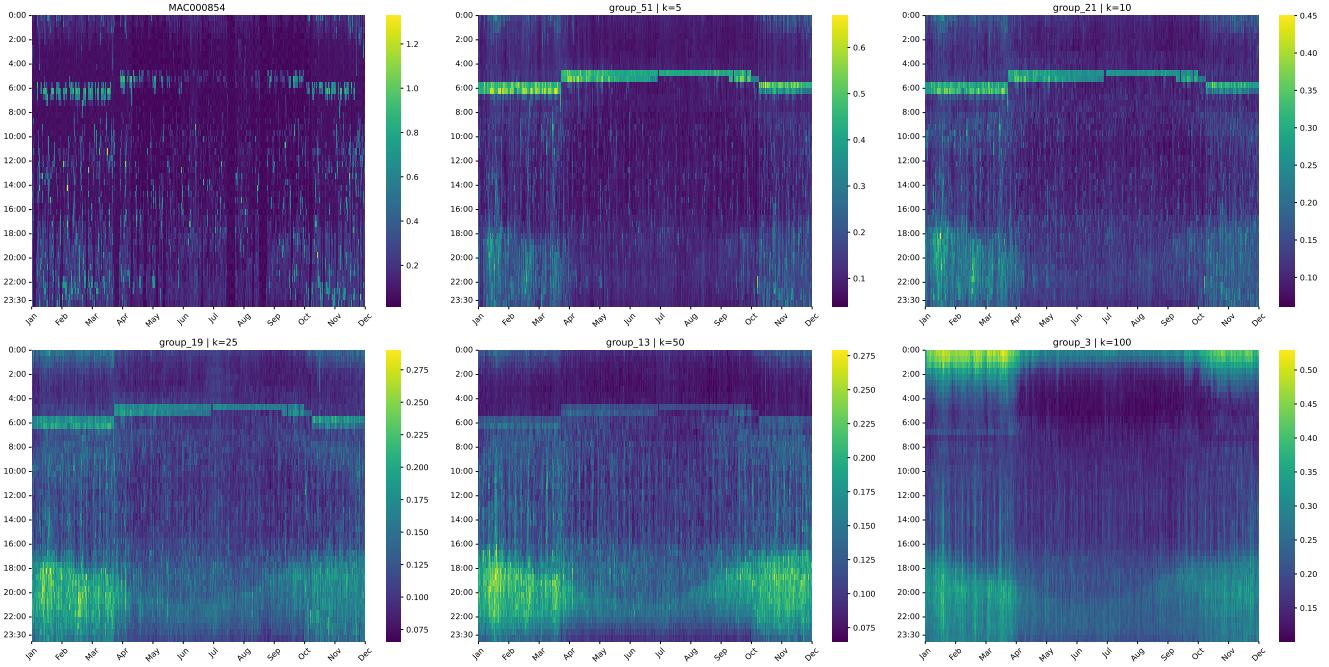


Fig. 7. Visualization of a single household's annual load profile (subplot 1) and the corresponding correctly re-linked micro-aggregated group profiles across increasing levels of k -anonymity (subplots 2 - 6). Group labels correspond to different aggregation levels, ranging from $k = 5$ to $k = 100$. The case $k = 200$ is omitted for clarity, as it closely resembles the $k = 100$ case.

associated with cyclic periods or specific behavioral routines, are sufficiently distinctive to persist through aggregation.

In contrast, in Appendix C Figure 8 presents a counterexample in which the target household exhibits typical seasonal heating behavior, characterized by higher winter energy usage, but lacks a unique fine-grained structure. This pattern is not sufficiently specific to enable successful linkage. As the aggregation level increases, the household's signal becomes diluted within the group, underscoring the weakening of identifiability for more typical load profiles.

Overall, our results highlight the inherent difficulty of anonymizing load profiles while maintaining utility. They further illustrate that the same properties which enable microaggregation to preserve forecasting accuracy—such as temporal coherence and low within-group variance—can also weaken its resistance to inference attacks. This tension between utility and privacy underscores the need for careful evaluation when applying microaggregation in practice.

From a practical standpoint, our findings suggest that the privacy-utility trade-off is not static; rather, it varies depending on the intended application. For example, when energy suppliers or aggregators aim to generate forecasts at an aggregated level, higher degrees of microaggregation can be beneficial. By reducing data variability, microaggregation simplifies the forecasting task, improves pattern recognition, and can lead to lower forecasting errors, ultimately reducing supplier imbalance costs.

Conversely, when the goal is to produce individual-level forecasts, microaggregation, while enhancing privacy protections—significantly hampers predictive accuracy and might lead to higher costs. Therefore, the utility of microaggregation is task-dependent, making it inappropriate to prescribe optimal levels based solely on our current study.

Nonetheless, our utility-privacy trade-off exploration provides a solid first step for the discussion among energy suppliers, aggregators, and even individual consumers acting as data controllers. It encourages them to consider the emerging potential of data markets and the possibility of sharing their data while ensuring anonymization. Once personal data is micro-aggregated, it no longer qualifies as personally identifiable, allowing third parties to derive meaningful utility while still preserving privacy for the consumer and adhering to regulation.

Consequently, while this study provides an initial evaluation of utility and privacy trade-offs in micro-aggregated smart meter data, several limitations must be acknowledged:

(i) *Closed-World Assumption*: Our exploration is conducted under a closed-world scenario in which every pseudonymized household in the attacker’s dataset corresponds exactly to one group in the micro-aggregated dataset. This setup simplifies the group membership inference task and represents a best-case condition for the attacker. However, in real-world deployments, datasets may contain unknown or unmatched households, and aggregation schemes may involve partial group overlap or missing data. Future work should consider open-world scenarios to more accurately capture practical attack feasibility and generalizability.

(ii) *Dataset scope*: The study relies exclusively on the LLC dataset. While the dataset’s resolution and socio-demographic richness make it ideal for controlled experimentation, it may not reflect regional or temporal variations in consumption patterns across different populations or grid infrastructures. Additional evaluations across diverse smart meter datasets would be necessary to validate the robustness of our findings and extend their applicability.

(iii) *Trade-off visualization and metric integration*: Our utility-privacy trade-off exploration uses a dual-axis plot to jointly visualize forecasting error (MAE) and inference success (RASR) across aggregation levels. While this offers an intuitive comparison, it lacks a unified metric framework and does not account for the differing scales or units of the two axes. As a result, visual interpretation alone may overstate or understate the trade-off severity. More principled methods from rate-distortion theory, such as Lagrangian optimization or Pareto efficiency frontiers, could help formalize trade-off quantification and support more rigorous privacy-preserving data publishing decisions.

9 CONCLUSION AND OUTLOOK

This work has introduced a replicable research methodology to systematically explore the utility-privacy trade-off in the context of smart meter data sharing. Focusing on microaggregation as an anonymization mechanism, we evaluated its dual impact on household-level load forecasting and privacy risks from group membership inference.

To this end, we formulated an unsupervised, similarity-based group membership inference attack, incorporating a novel two-stage majority voting scheme to re-link pseudonymized household profiles to their anonymized groups. In contrast to prior work focused on binary group membership decisions, our more complex matching-based formulation exposes higher privacy risks, particularly when combined with advanced feature representations.

Our experiments reveal that structural patterns in energy consumption persist even after anonymization via MDAV. These residual patterns can still be exploited by both domain-informed features and deep learning models. Convolutional autoencoders, in particular, demonstrate strong performance in inferring group membership at low anonymity levels (e.g., $k = 5$). At the same time, these patterns enable accurate household-level forecasting even under higher anonymization (e.g., $k = 100$). This underscores the ongoing tension between maintaining data utility and ensuring strong privacy protection.

Future work could extend this methodology in several directions. First, evaluations could be broadened to include multiple datasets and more relaxed assumptions, such as open-world scenarios involving previously unseen households. Second, the utility-privacy trade-off might be more rigorously formalized as a multi-objective optimization (MOO) problem, inspired by rate-distortion theory. Methods such as Pareto frontier analysis or Lagrange optimization could then be employed to identify optimal operating points that balance privacy protection and analytical utility.

In addition, future research might explore alternative microaggregation schemes or enhance representation learning through more expressive models. For instance, denoising autoencoders trained to reconstruct household-level signals from aggregated inputs could inadvertently learn to suppress inter-household variance, thereby enabling re-identification. These models could introduce new privacy risks, underscoring the need for careful assessment of anonymization guarantees in time-series data.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the authors used DeepL [19], ChatGPT [53] and Grammarly [40] to paraphrase and fix grammatical mistakes. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

ACKNOWLEDGMENTS

This research was funded in part by the Luxembourg National Research Fund (FNR) and PayPal, PEARL grant reference 13342933/Gilbert Fridgen, by FNR grant reference HPC BRIDGES/2022_Phase2/17886330/DELPHI. For the purpose of open

access and fulfilling the obligations arising from the grant agreement, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission. Additionally, this paper has been supported by Enovos. Funding from the Federal State of Salzburg through project TRAMPOLIN-IT is gratefully acknowledged.

REFERENCES

- [1] Joana M Abreu, Francisco Câmara Pereira, and Paulo Ferrão. 2012. Using Pattern Recognition to Identify Habitual Behavior in Residential Electricity Consumption. *Energy and Buildings* 49 (2012), 479–487. <https://doi.org/10.1016/j.enbuild.2012.02.044>
- [2] Kayode S Adewole and Vicenç Torra. 2022. DFTMicroagg: a dual-level anonymization algorithm for smart grid data. *International Journal of Information Security* 21, 6 (2022), 1299–1321.
- [3] Muhammad Rizwan Asghar, György Dán, Daniele Miorandi, and Imrich Chlamtac. 2017. Smart Meter Data Privacy: A Survey. *IEEE Communications Surveys and Tutorials* 19, 4 (2017), 2820–2835. <https://doi.org/10.1109/COMST.2017.2720195>
- [4] Christian Beckel, Wilhelm Kleiinger, Romano Cicchetti, Thorsten Staake, and Silvia Santini. 2014. The ECO Data Set and the Performance of Non-Intrusive Load Monitoring Algorithms. In Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings (Memphis, Tennessee). *Proceedings of the 1st ACM International Conference on Embedded Systems for Energy-Efficient Buildings*, 80–89. <https://doi.org/10.1145/2674061.2674064>
- [5] Christian Beckel, Leyna Sadamori, and Silvia Santini. 2013. Automatic socio-economic classification of households using electricity consumption data. *e-Energy 2013 - Proceedings of the 4th ACM International Conference on Future Energy Systems* (2013), 75–86. <https://doi.org/10.1145/2487166.2487175>
- [6] Christian Beckel, Heinz Serfas, Elmar Zeeb, Guido Moritz, Frank Golatowski, and Dirk Timmermann. 2011. Requirements for Smart Home Applications and Realization with WS4D-PipesBox, In ETFA2011 (Toulouse). *Proceedings of the 16th IEEE Conference on Emerging Technologies and Factory Automation (ETFA 2011)*, 1–8. <https://doi.org/10.1109/ETFA.2011.6059229>
- [7] Jens-Matthias Bohli, Christoph Sorge, and Osman Ugus. 2010. A Privacy Model for Smart Metering. In *2010 IEEE International Conference on Communications Workshops*. 1–5. <https://doi.org/10.1109/ICCW.2010.5503916>
- [8] Niklas Buescher, Spyros Boukros, Stefan Bauregger, and Stefan Katzenbeisser. 2017. Two is not enough: Privacy assessment of aggregation schemes in smart metering. *Proceedings on Privacy Enhancing Technologies* (2017).
- [9] Cristian Challu, Kim G Olivares, Boris N Oreshkin, Federico Garza Ramirez, Max Mergenthaler Canseco, and Artur Dubrawski. 2023. Nhts: Neural hierarchical interpolation for time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 6989–6997.
- [10] Dong Chen, Sean Barker, Adarsh Subbaswamy, David Irwin, and Prashant Shenoy. 2013. Non-Intrusive Occupancy Monitoring using Smart Meters. In Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings. *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings - BuildSys'13*, 1–8. <https://doi.org/10.1145/2528282.2528294>
- [11] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA). Association for Computing Machinery, New York, NY, USA, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [12] Ying Chen, Peter B. Luh, Che Guan, Yige Zhao, Laurent D. Michel, Matthew A. Coolbeth, Peter B. Friedland, and Stephen J. Rourke. 2010. Short-Term Load Forecasting: Similar Day-Based Wavelet Neural Networks. *IEEE Transactions on Power Systems* 25, 1 (2010), 322–330. <https://doi.org/10.1109/TPWRS.2009.2030426>
- [13] Sara Clempunt, Mustafa A. Mustafa, Eduard Marin, and Bart Prenel. 2018. De-pseudonymization of Smart Metering Data: Analysis and Countermeasures. In *2018 Global Internet of Things Summit (GloTS)*. 1–6. <https://doi.org/10.1109/GIOTS.2018.8534430>
- [14] European Commission. 1995. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Communities* L 281 (23 Nov. 1995), 31–50. <https://eur-lex.europa.eu/eli/dir/1995/46/eng> Accessed May 8, 2025.
- [15] Creos Luxembourg S.A. 2024. *Annual Report 2023*. Technical Report. Creos Luxembourg S.A. https://www.creos-net.lu/fileadmin/dokumente/downloads/gb_creos_annual_report_2023.pdf?t=1746705710930 Accessed May 8, 2025.
- [16] Ana-Maria Cretu, Miruna Rusu, and Yves-Alexandre de Montjoye. 2024. Re-pseudonymization Strategies for Smart Meter Data Are Not Robust to Deep Learning Profiling Attacks. In *Proceedings of the Fourteenth ACM Conference on Data and Application Security and Privacy* (Porto, Portugal) (CODASPY '24). Association for Computing Machinery, New York, NY, USA, 295–306. <https://doi.org/10.1145/3626232.3653272>
- [17] Jean-Michel D. 2019. Smart meter data from London area. <https://www.kaggle.com/jeannimidev/smart-meters-in-london>
- [18] Yves-Alexandre de Montjoye, Cesar A Hidalgo, Michel Verleysen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3 (2013), 1376.
- [19] DeepL SE. 2025. DeepL. <https://www.deepl.com/write>.
- [20] Joaquin Delgado Fernandez, Sergio Potenciano Menci, and Alessio Magitteri. 2025. Forecasting Anonymized Electricity Load Profiles. *arXiv:2501.06237 [cs.CR]* <https://arxiv.org/abs/2501.06237>
- [21] Aljoscha Dietrich, Dominik Leibenger, and Christoph Sorge. 2020. On the lack of anonymity of anonymized smart meter data: An empiric study. In *2020 IEEE 45th Conference on Local Computer Networks (LCN)*. IEEE, 405–408.
- [22] Josep Domingo-Ferrer, Josep Maria Mateo-Sanz, and Antoni Torres. 2004. *Privacy in Statistical Databases: Microaggregation and Perturbation Techniques*. Springer.
- [23] Josep Domingo-Ferrer and Vicenç Torra. 2005. Ordinal, continuous and heterogeneous k-anonymity through microaggregation. *Data Mining and Knowledge Discovery* 11 (2005), 195–212.
- [24] Josep Domingo-Ferrer and Vicenç Torra. 2008. A critique of k-anonymity and some of its enhancements. In *Third International Conference on Availability, Reliability and Security*. IEEE, 990–993.
- [25] Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
- [26] EDDIE Project Consortium. 2025. EDDIE – European Distributed Data Infrastructure for Energy. <https://eddie.energy/> Accessed: 2025-05-09.
- [27] Eurelectric. 2024. *A snapshot of Cybersecurity in the EU*. Technical Report. Eurelectric. <https://www.eurelectric.org/wp-content/uploads/2024/11/A-Eurelectric-snapshot-of-Cybersecurity-2024-11-18-FINAL.pdf> Accessed: 2025-05-09.
- [28] European Commission. 2012. Recommendation of 9 March 2012 on preparations for the roll-out of smart metering systems. *Official Journal of the European Union*. Available online: <https://op.europa.eu/en/publication-detail/-/publication/a5daa8c6-8f11-4e5e-9634-3f224af571a6/language-en> [2025-04-03].
- [29] European Commission. 2014. *Cost-benefit analyses and state of play of smart metering deployment in the EU-27*. SWD(2014) 189 final. European Commission. Available online: <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=140308459559&uri=SWD:2014:189:FIN> [2025-04-03].
- [30] European Commission. 2022. Best Practice for Energy Data Sharing. https://commission.europa.eu/document/download/0518ac9e-faa3-472a-9a8e-1fa6a87ccc38_en?filename=final_agenda_1602.pdf Accessed: 2025-05-09.
- [31] European Commission. 2025. Common European Data Spaces. <https://digital-strategy.ec.europa.eu/en/policies/data-spaces> Accessed: 2025-05-09.
- [32] European Commission. 2025. The Digital Services Act. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services_act_en Accessed: 2025-05-09.
- [33] European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* L 119 (May 2016), 1–88. <https://eur-lex.europa.eu/eli/reg/2016/679/oj> Published on 4 May 2016.
- [34] European Union. 2022. Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act). *Official Journal of the European Union* L 152/1 (May 2022). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R0868> Accessed: 2025-05-09.
- [35] European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). *Official Journal of the European Union* L 204/1 (12 July 2024), 1–124. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689> Accessed: 2025-05-09.
- [36] Zhong Fan, Parag Kulkarni, Sedat Gormus, Costas Eftymiou, Georgios Kalogridis, Mahesh Sooriyabandara, Ziming Zhu, Sangarapillai Lambotharan, and Woon Hau Chin. 2013. Smart grid communications: Overview of research challenges, solutions, and standardization activities. *IEEE Communications Surveys and Tutorials* 15 (2013), 21–38. Issue 1. <https://doi.org/10.1109/SURV.2011.122211.00021>
- [37] Cornelia Ferner, Günther Eibl, Andreas Unterweger, Sebastian Burkhart, and Stefan Wegenkittl. 2019. Pool Detection from Smart Metering Data with Convolutional Neural Networks. *Energy Informatics* 2 (2019), 1–9. <https://doi.org/10.1186/s42162-019-0097-8>

[38] Sören Finster and Ingmar Baumgart. 2015. Privacy-Aware Smart Metering: A Survey. *IEEE Communications Surveys and Tutorials* 17, 2 (2015), 1088–1101. <https://doi.org/10.1109/COMST.2015.2425958>

[39] Benjamin CM Fung, Ke Wang, Rui Chen, and Philip S Yu. 2010. Privacy-preserving data publishing: A survey of recent developments. *ACM Computing Surveys (CSUR)* 42, 4 (2010), 1–53.

[40] Grammarly, Inc. 2025. Grammarly: AI Writing and Grammar Checker. <https://www.grammarly.com>. Accessed: 2025-05-16.

[41] Konstantin Hopf, Mariya Sodenkamp, Ilya Kozlovkiy, and Thorsten Staake. 2016. Feature extraction and filtering for household classification based on smart electricity meter data. *Computer Science - Research and Development* 31 (8 2016), 141–148. Issue 3. <https://doi.org/10.1007/S00450-014-0294-4>

[42] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S. Yu, and Xuyun Zhang. 2022. Membership Inference Attacks on Machine Learning: A Survey. *ACM Comput. Surv.* 54, 11, Article 235 (2022), 37 pages. <https://doi.org/10.1145/3523273>

[43] Marek Jawurek, Martin Johns, and Konrad Rieck. 2011. Smart metering de-pseudonymization. In *Proceedings of the 27th Annual Computer Security Applications Conference* (Orlando, Florida, USA) (ACSAC '11). Association for Computing Machinery, New York, NY, USA, 227–236. <https://doi.org/10.1145/2076732.2076764>

[44] Shubhankar Kapoor, Björn Sturmberg, and Marnie Shaw. 2020. A review of publicly available energy data sets. *Watwatchers' My Energy Marketplace (MEM) (The Australian National University, p. 2020. Canberra, Australia 2020)*.

[45] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: a highly efficient gradient boosting decision tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 3149–3157.

[46] Younghun Kim, Edith C-H Ngai, and Mani B Srivastava. 2011. Cooperative state estimation for preserving privacy of user behaviors in smart grid. In *2011 IEEE international conference on smart grid communications (SmartGridComm)*. IEEE, 178–183.

[47] Wilhelm Kleiminger, Christian Beckel, Thorsten Staake, and Silvia Santini. 2013. Occupancy Detection from Electricity Consumption Data. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys '13)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2528282.2528295>

[48] J Zico Kolter and Tommi Jaakkola. 2012. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. *Journal of Machine Learning Research - Proceedings Track* 22 (4 2012), 1472–1482.

[49] Mikhail A Lisovich and Stephen B Wicker. 2008. Privacy Concerns in Upcoming Residential and Commercial Demand-Response Systems. *IEEE Proceedings on Power Systems* 1, 1 (2008), 1–10.

[50] Daniel L. Marino, Kasun Amarasinghe, and Milos Manic. 2016. Building energy load forecasting using Deep Neural Networks. In *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*. 7046–7051. <https://doi.org/10.1109/IECON.2016.7793413>

[51] Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. 2010. Private memoirs of a smart meter, In Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building (New York, NY, USA). *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, 61–66.

[52] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 111–125.

[53] Inc. OpenAI. 2025. ChatGPT. <https://chatgpt.com/>.

[54] Boris N Oreshkin, Dmitri Carpow, Nicolas Chapados, and Yoshua Bengio. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437* (2019).

[55] Fotios Petropoulos, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J Bessa, Jakub Bijak, John E Boylan, et al. 2022. Forecasting: theory and practice. *International Journal of Forecasting* 38, 3 (2022), 705–871.

[56] Dejan Radovanovic, Maximilian Schiril, Andreas Unterweger, and Günther Eibl. 2025. Predicting Socio-Demographic Characteristics from Load Profiles with Varying Time Granularities. In *Proceedings of the 14th International Conference on Smart Cities and Green ICT Systems - SMARTGREENS*. INSTICC, SciTePress, 87–98. <https://doi.org/10.5220/0013217400003953>

[57] Dejan Radovanovic, Andreas Unterweger, Günther Eibl, Dominik Engel, and Johannes Reichl. 2022. How unique is weekly smart meter data? *Energy Informatics* 5 (2022), 1–13. <https://doi.org/10.1186/s42162-022-00205-8>

[58] Rouzbeh Razavi, Amin Gharipour, Martin Fleury, and Ikpe Justice Akpan. 2019. Occupancy detection of residential buildings using smart meter data: A large-scale study. *Energy and Buildings* 183 (2019), 195–208. <https://doi.org/10.1016/j.enbuild.2018.11.025>

[59] Claudio Sarra. 2024. Artificial Intelligence in Decision-making: A Test of Consistency between the "EU AI Act" and the "General Data Protection Regulation". *Athens Journal of Law* 11 (10 2024), 1–17. <https://doi.org/10.30958/ajl.11-1-3>

[60] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership Inference Attacks against Machine Learning Models. *arXiv:1610.05820 [cs.CR]* <https://arxiv.org/abs/1610.05820>

[61] Evangelos Spiliotis, Konstantinos Nikolopoulos, and Vassilios Assimakopoulos. 2019. Tales from tails: On the empirical distributions of forecasting errors and their implication to risk. *International Journal of Forecasting* 35, 2 (2019), 687–698. <https://doi.org/10.1016/j.ijforecast.2018.10.004>

[62] Statista. 2023. Number of Smart Meters in the United States. <https://www.statista.com/statistics/499704/number-of-smart-meters-in-the-united-states/>. Accessed: 2025-04-14.

[63] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.

[64] Matthias Templ, Alexander Kowarik, and Bernhard Meindl. 2015. Statistical Disclosure Control for Micro-Data Using the R Package sdcMicro. *Journal of Statistical Software* 67, 4 (2015), 1–36. <https://doi.org/10.18637/jss.v067.i04>

[65] Simon Tindemans. 2023. Low Carbon London smart meter data (refactored). <https://doi.org/10.4121/fbb775b-48d8-469f-a39b-b64488bfddfd.v1>

[66] UK Government. 2019. Statistical Release and Data: Smart Meters Great Britain, Quarter 3 2019. <https://www.gov.uk/government/statistics/statistical-release-and-data-smart-meters-great-britain-quarter-3-2019>. Accessed: 2025-04-14.

[67] Andreas Unterweger, Sanaz Taheri-Boshrooyeh, Günther Eibl, Fabian Knirsch, Alptekin Küpcü, and Dominik Engel. 2019. Understanding Game-Based Privacy Proofs for Energy Consumption Aggregation Protocols. , 5514–5523 pages. <https://doi.org/10.1109/TSG.2018.2883951>

[68] Yi Wang, Qixin Chen, Dahua Gan, Jingwei Yang, Daniel S. Kirschen, and Chongqing Kang. 2019. Deep learning-based socio-demographic information identification from smart meter data. *IEEE Transactions on Smart Grid* 10 (5 2019), 2593–2602. Issue 3. <https://doi.org/10.1109/TSG.2018.2805723>

[69] Yi Wang, Qixin Chen, Tao Hong, and Chongqing Kang. 2019. Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges. *IEEE Transactions on Smart Grid* 10, 3 (2019), 3125–3148. <https://doi.org/10.1109/TSG.2018.2818167>

[70] WP216. 2014. *Opinion 05/2014 on Anonymisation Techniques*. Technical Report WP216, Article 29 Data Protection Working Party. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf Adopted on 10 April 2014.

[71] Jing Xiong, Pengyang Zhou, Alan Chen, and Yu Zhang. 2021. Attention-based Neural Load Forecasting: A Dynamic Feature Selection Approach. In *2021 IEEE Power Energy Society General Meeting (PESGM)*. IEEE, 01–05. <https://doi.org/10.1109/pesgm46819.2021.9637992>

[72] Shiliang Zhang, Sabita Maharjan, Lee Andrew Bygrave, and Shui Yu. 2025. Data Sharing, Privacy and Security Considerations in the Energy Sector: A Review from Technical Landscape to Regulatory Specifications. *arXiv:2503.03539 [cs.CR]*

[73] Ioannis Zografopoulos, Juan Ospina, Xiaorui Liu, and Charalambos Konstantinou. 2021. Cyber-Physical Energy Systems Security: Threat Modeling, Risk Assessment, Resources, Metrics, and Case Studies. *IEEE Access* 9 (2021), 29775–29818. <https://doi.org/10.1109/ACCESS.2021.3058403>

A COMPLETE FORECASTING PERFORMANCE

Table 4 displays the forecasting results across various anonymity levels, ranging from $k = 5$ to $k = 200$. It also includes performance evaluations based on four key metrics: MAE, MSE, RMSE, and SMAPE.

Table 4. Evaluation metrics by model and anonymity regime (sorted by increasing anonymity).

| Model | Anonymity | MAE | MSE | RMSE | SMAPE |
|-------------------|-----------|-------|-------|-------|-------|
| Linear Regression | k=5 | 0.119 | 0.052 | 0.178 | 0.239 |
| | k=10 | 0.120 | 0.052 | 0.179 | 0.244 |
| | k=25 | 0.122 | 0.053 | 0.181 | 0.257 |
| | k=50 | 0.126 | 0.055 | 0.185 | 0.273 |
| | k=100 | 0.130 | 0.057 | 0.189 | 0.289 |
| | k=200 | 0.135 | 0.061 | 0.195 | 0.307 |
| LGBM | k=5 | 0.109 | 0.052 | 0.175 | 0.216 |
| | k=10 | 0.111 | 0.052 | 0.175 | 0.219 |
| | k=25 | 0.114 | 0.053 | 0.178 | 0.225 |
| | k=50 | 0.117 | 0.055 | 0.182 | 0.229 |
| | k=100 | 0.120 | 0.057 | 0.186 | 0.235 |
| | k=200 | 0.124 | 0.061 | 0.193 | 0.243 |
| XGBoost | k=5 | 0.115 | 0.050 | 0.174 | 0.232 |
| | k=10 | 0.116 | 0.051 | 0.176 | 0.233 |
| | k=25 | 0.118 | 0.053 | 0.179 | 0.234 |
| | k=50 | 0.121 | 0.055 | 0.183 | 0.238 |
| | k=100 | 0.124 | 0.058 | 0.188 | 0.245 |
| | k=200 | 0.127 | 0.062 | 0.194 | 0.254 |
| NBEATS | k=5 | 0.114 | 0.052 | 0.174 | 0.239 |
| | k=10 | 0.116 | 0.052 | 0.175 | 0.245 |
| | k=25 | 0.121 | 0.054 | 0.179 | 0.257 |
| | k=50 | 0.123 | 0.056 | 0.180 | 0.262 |
| | k=100 | 0.129 | 0.060 | 0.185 | 0.273 |
| | k=200 | 0.137 | 0.065 | 0.192 | 0.288 |
| NHITS | k=5 | 0.114 | 0.052 | 0.173 | 0.236 |
| | k=10 | 0.117 | 0.053 | 0.175 | 0.244 |
| | k=25 | 0.124 | 0.055 | 0.179 | 0.259 |
| | k=50 | 0.129 | 0.060 | 0.184 | 0.274 |
| | k=100 | 0.136 | 0.064 | 0.190 | 0.291 |
| | k=200 | 0.139 | 0.068 | 0.194 | 0.298 |

B ATTACKER ARCHITECTURES

Table 5 showcases the architecture used for the Convolutional autoencoder while Table 6 on the other hand contains the architecture of Recurrent Autoencoder use for the attacker model.

Table 5. Architecture of the Convolutional Autoencoder.

| Layer | Type | Output Shape |
|----------------|-------------------------------|--------------|
| <i>Encoder</i> | | |
| Input | Input Layer | (7, 48, 1) |
| Conv2D | 64 filters, (7,8), ReLU, same | (7, 48, 64) |
| MaxPooling2D | Pool size (1,2) | (7, 24, 64) |
| Conv2D | 32 filters, (4,4), ReLU, same | (7, 24, 32) |
| MaxPooling2D | Pool size (1,2) | (7, 12, 32) |
| Conv2D | 16 filters, (2,2), ReLU, same | (7, 12, 16) |
| MaxPooling2D | Pool size (1,2) | (7, 6, 16) |
| Conv2D | 16 filters, (7,2), ReLU, same | (7, 6, 16) |
| Conv2D | 8 filters, (7,2), ReLU, same | (7, 6, 8) |
| MaxPooling2D | Pool size (7,1) | (1, 6, 8) |
| Flatten | - | (48) |
| <i>Decoder</i> | | |
| Reshape | (1, 6, 8) | (1, 6, 8) |
| UpSampling2D | Size (7,1) | (7, 6, 8) |
| Conv2D | 8 filters, (7,2), ReLU, same | (7, 6, 8) |
| Conv2D | 16 filters, (7,2), ReLU, same | (7, 6, 16) |
| UpSampling2D | Size (1,2) | (7, 12, 16) |
| Conv2D | 16 filters, (2,2), ReLU, same | (7, 12, 16) |
| UpSampling2D | Size (1,2) | (7, 24, 16) |
| Conv2D | 32 filters, (4,4), ReLU, same | (7, 24, 32) |
| UpSampling2D | Size (1,2) | (7, 48, 32) |
| Conv2D | 64 filters, (7,8), ReLU, same | (7, 48, 64) |
| Conv2D | 1 filter, (7,8), Linear, same | (7, 48, 1) |

Table 6. Architecture of the Recurrent Autoencoder.

| Layer | Type / Parameters | Output Shape |
|-----------------|-------------------------------------|--------------|
| <i>Encoder</i> | | |
| Input | Input Layer (7, 48) | (7, 48) |
| LSTM | 64 units, ReLU, return_seq | (7, 64) |
| Dropout | rate = 0.2 | (7, 64) |
| LSTM | 64 units, ReLU, return_seq | (7, 64) |
| Dropout | rate = 0.2 | (7, 64) |
| LSTM | 32 units, ReLU L2 reg. on kernel | (32) |
| <i>Decoder</i> | | |
| Input | Input Layer (32) | (32) |
| RepeatVector | Repeats to 7 time steps | (7, 32) |
| LSTM | 32 units, ReLU, return_seq | (7, 32) |
| LSTM | 64 units, ReLU, return_seq | (7, 64) |
| LSTM | 64 units, ReLU, return_seq | (7, 64) |
| TimeDistributed | Dense(48) per time step | (7, 48) |

C HOUSEHOLD AND GROUP VISUALIZATION

Figure 8 contains the visualization of a single household's annual load profile (subplot 1) and the wrongly predicted micro-aggregated group profiles across increasing levels of k -anonymity (subplots 2 - 6). Group labels correspond to different aggregation levels, ranging from $k = 5$ to $k = 100$. The case $k = 200$ is omitted for clarity, as it closely resembles the $k = 100$ case.

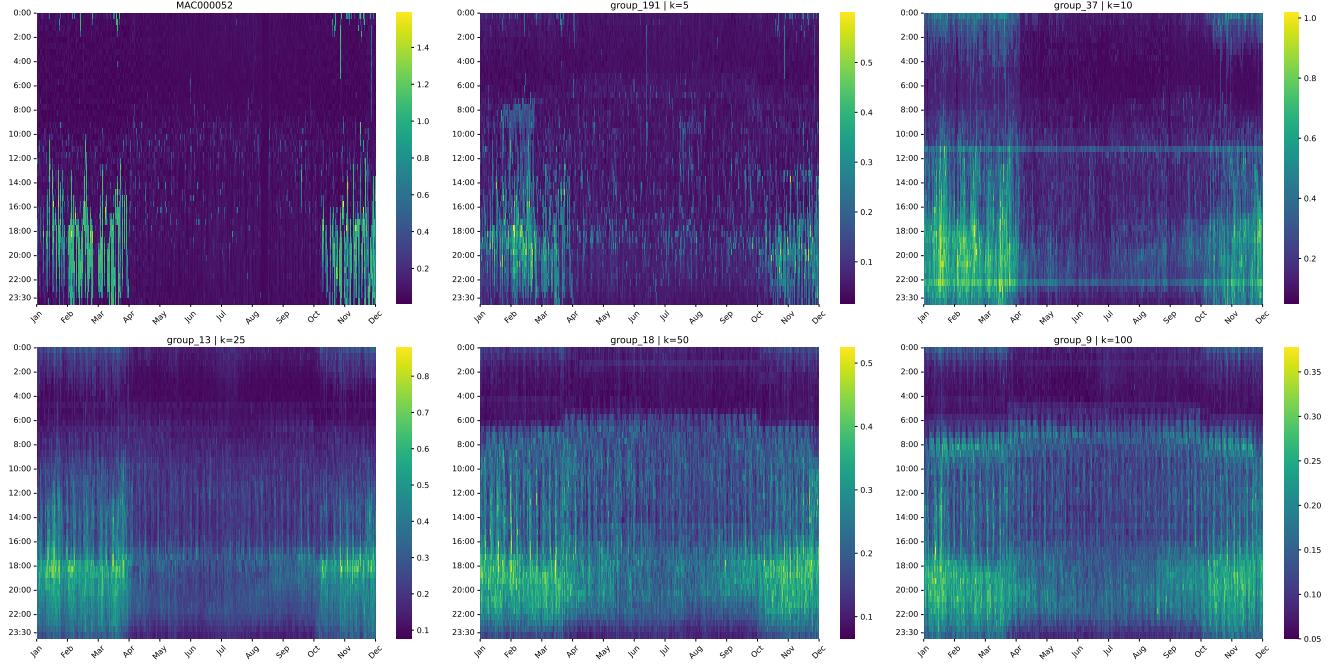


Fig. 8. Household and group visualization.

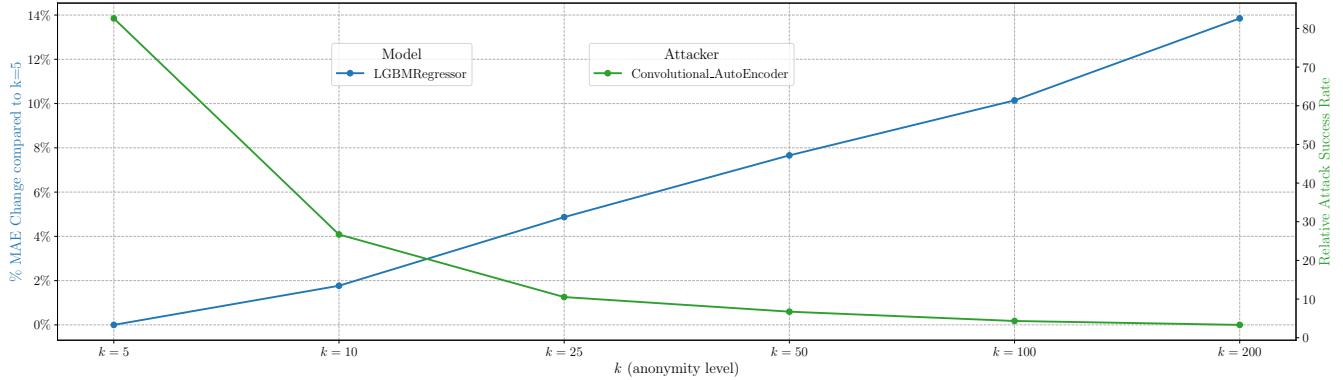


Fig. 9. Visualization of MAE degradation against RASR.

D TRADE-OFF BETWEEN PRIVACY DEGRADATION AND RASR

Figure 9 visualizes the trade-off between utility and privacy for different levels of k -anonymity. Utility is quantified by the mean absolute degradation in forecasting error (% kWh, left y-axis, shown in blue) using a LGBM Regressor, while privacy is measured by the RASR (right y-axis, shown in green).