

A blind source separation framework to monitor sectoral power demand from grid-scale load measurements

Guillaume Koechlin^{a,*}, Filippo Bovera^b, Elena Degli Innocenti^c, Barbara Santini^c, Alessandro Venturi^c, Simona Vazio^c, Piercesare Secchi^a

^a MOX, Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci 32, Milano, 20133, MI, Italy

^b Department of Energy, Politecnico di Milano, Via Lambruschini 4a, Milano, 20156, MI, Italy

^c Statistics Group, Terna S.p.A, Viale Egidio Galbani, 70, Rome, 00156, RM, Italy

ARTICLE INFO

Keywords:

Load disaggregation
Blind source separation
Load profile
Non-intrusive load monitoring
Non-negative matrix factorization

ABSTRACT

As demand-side flexibility becomes increasingly necessary to integrate variable renewable energy, understanding electricity demand composition across different grid levels is essential. However, at regional and national scales, visibility into the relative contributions of different consumer categories remains limited due to the complexity and cost of collecting end-use consumption data. To address this challenge, we propose a blind source separation framework to disaggregate open-access high-voltage grid load measurements into sectoral contributions. The approach relies on a constrained variant of non-negative matrix factorization, termed linearly-constrained non-negative matrix factorization (LCNMF), which allows prior information to be incorporated as linear constraints on the factor matrices, thereby providing weak supervision of the separation process. The framework is evaluated using Italian national load data from 2021 to 2023. Results demonstrate the identifiability of residential, services, and industrial load components and provide monthly sectoral consumption estimates consistent with reported statistics. The proposed method is generalizable and applicable to load disaggregation problems across multiple grid scales where disaggregated measurements are unavailable.

1. Introduction

1.1. Motivations

Monitoring and understanding electricity consumption dynamics is essential for power system operation. At the national level, electricity consumption (and generation) is typically monitored by the transmission system operator (TSO) both at the level of geographical zones, such as states, regions or power zones, and across different sectors often grouped into three or four categories: industry, commercial & public services (or simply *commercial* or *services*), residential (or *households*) and sometimes agriculture [1,2]. While zone-level consumption is directly accessible to the TSO at high temporal resolution via measurements on the high-voltage network, sector-level consumption is considerably more difficult to monitor. It must be retrieved from the medium- and low-voltage networks,¹ which fall outside the TSO's scope and are managed by numerous distribution system operators (DSOs). This process thus involves

aggregating numerous building-level load measurements, which are not centrally collected. As a result, sector-specific consumption data are often available only at very poor temporal resolution, mostly on an annual basis, and with significant reporting delays [3].

However, estimating sector-specific consumption at finer temporal resolutions, such as monthly or daily, or even directly monitoring the load at the sector level, is of particular interest to TSOs and other power system actors for two reasons:

First, it obviously provides valuable knowledge on sectors themselves. For instance, global sectors' typical load profiles and granular consumption can be used for the calibration and validation of building stock energy models, especially in top-down approaches [4]. It is also valuable for demand-side management as it allows an understanding of load composition during peak hours—which may vary depending on the type of day or the season—and therefore helps in designing efficient sector-specific demand response programs [5]. Real-time consumption

* Corresponding author.

Email address: guillaume.koechlin@polimi.it (G. Koechlin).

¹ With the exception of a few industrial consumers directly connected to the high-voltage grid.

monitoring can also help rapidly detect and quantify how the different sectors respond to economic shocks (e.g., energy price spikes) [6,7], policy interventions (e.g., lockdowns) [8–10] or extreme weather events (e.g., heat waves) [11], which is increasingly important in a context of rising demand-side and generation flexibility needs. Past and current efforts on sectors' power consumption monitoring rely on small consumer samples, which are not always representative nor accessible [3,12], and are only available a posteriori. Finally, sector-level demand at higher frequency is obviously not only useful to the power sector but also to economic institutions and policymakers to monitor national economic activity [13] with greater precision.

Second, it can enhance aggregate load modeling and forecasting [14–16]. A significant portion of the variability in the national load curve arises from the dynamic contributions of different sectors [15]. For example, weekday load profiles are typically flatter than those on weekends due to a stronger industry contribution, whose consumption pattern is rather stable throughout the day. It may also improve the modeling of weather or calendar effects on the aggregate load curve, as different load components (i.e., sectors) may respond differently to specific conditions. For instance, the household and services sectors may show higher temperature sensitivity in summer due to widespread air conditioning usage, compared to the industrial sector [17]. Additionally, the long-term evolution of sector-specific load profiles may be easier to model and project in scenarios—especially in the context of electrification—than the aggregate national load [6,14,18].

1.2. Objectives of the work

In this work, we demonstrate that, despite the absence of actual measurements, it is possible to monitor the hourly consumption of the residential, commercial and industrial sectors in real-time from the aggregate open-access TSO load measurements, by applying signal separation techniques. The signal separation framework we propose enables both the extraction of sector-specific daily load profiles and the now-casting of daily or monthly sectoral consumption. We formulate this unsupervised load disaggregation task as a *blind* source separation (BSS) problem, showing that the incorporation of sector-level information usually available—such as annual demand and monthly consumption proxies—as boundary conditions allows us to identify characteristic sectors' load profiles and obtain an accurate sector-wise disaggregation. The BSS problem is formalized as a constrained variant of non-negative matrix factorization (NMF) [19], which is termed linearly-constrained non-negative matrix factorization (LCNMF), and solved using a modified version of the multiplicative updates algorithm [20]. We apply this method to Italian national load data from 2021–2022 and show that the identified sector profiles for this period allow us to accurately disaggregate in real-time new load curves for which no sector-level information is available, using 2023 data for validation.

1.3. Paper structure

The remainder of this paper is structured as follows. Section 2 provides a detailed literature review and explains our contributions. Section 3 details the methodology, beginning with a description of the data sources (Section 3.1) and a qualitative introduction to the problem (Section 3.2) before presenting the mathematical and computational details (Sections 3.3 and 3.4). Finally, Section 4 presents the results, followed by a discussion in Section 5 and concluding remarks in Section 6.

2. Literature review

Statistical modeling and analysis of load curves is a major topic in power systems. A literature overview suggests that most works belong to one of the three following categories:

1. **Load forecasting:** The problem is to predict future load at various network scales from past measurements and influencing factors

such as weather and calendar effects. Though the results of the present work may be relevant for such purposes, the problem we solve is not a forecasting one and the corresponding literature is therefore not further presented here.

2. **Load profiling/classification:** This category concerns the analysis of typical patterns and similarities between load curves observed in different locations, for different day types, seasons or categories of consumers. The comparison with such works is relevant since our method involves identifying sector-specific load profiles.
3. **Load disaggregation:** We refer to load disaggregation as the task of separating unobserved devices, end-users or consumer groups contributions to observed aggregated load measurements. This is precisely the problem of interest. Since nearly all load disaggregation works are limited to device-level decomposition of buildings consumption, we selected those whose methodological framework is close to ours and this part of the review is essentially on the technical side.

In addition, we demonstrate that this work brings a genuine methodological novelty by extending the discussion to all contributions that have proposed constrained non-negative matrix factorization methods, regardless of the application.

2.1. Load profiling and classification

In load profiling and classification, a distinction can be made according to whether load measurements or models correspond to end-users or an aggregated level (such as district or regional/national level).

At the end-user level, typical residential, commercial and industrial load profiles can be inferred by averaging individual consumers' load curves within each sector. This is typically done with a manual weekday/weekend, season, region and household characteristics stratification or in a more data-driven fashion using clustering techniques. As highlighted in Ref. [3], thanks to the widespread diffusion of smart meters, the residential sector has received the most attention in the literature [21–26], while fewer works have studied commercial and industrial consumers [27,28]. Our study, instead, considers aggregate country-level profiles and their temporal variations rather than within-sector heterogeneity.

Regarding the aggregate level, most works use a so-called bottom-up approach [4]: The load profile of the aggregation level is built by aggregating load profiles of different consumer categories with weights reflecting the local consumer mix. For instance, Ref. [3] builds neighborhood, district and municipality-level load profiles, which are shown to be well-described by a three-level “residential”, “business”, and “mixed” classification. Their results show that at aggregate levels, only global sector categorization really makes the difference, rather than finer ones. Their study however lacks the inclusion of the industrial component, which is instead present in our work. Ref. [15] aggregates load curves of 4500 end-users in Denmark to derive national profiles for residential, commercial, public services, industry and agriculture sectors, finding strong weekly and seasonal variation except for industry. Later work [6] shows that sector profiles remain relatively stable over the 2019–2022 period. Finally, Refs. [14,18] adopt more of a top-down approach as they estimate the consumer mix in local areas in Denmark by combining an end-users load profiles dataset and transformer station aggregate measurements. Note that, similarly to Ref. [3], they rely on the assumption that load profiles of specific consumer categories are constant across locations and, differently from us, do not properly model sectors' load profile variations over time. Moving to the highest aggregation level, Ref. [29] proposed a methodology to build weather-dependent load profiles for European countries. They however do not address the weather-independent temporal variations of load profiles and sector-specific consumption patterns.

2.2. Load disaggregation

Prior research on load profiling and classification has consistently shown that aggregate curves from the substation to the national level contain identifiable sector components. However, while building-scale load disaggregation is a flourishing field of study, which we will discuss first, disaggregation at higher grid scales remains largely unexplored. We will review the few existing works on this topic in a subsequent section.

2.2.1. Non-intrusive load monitoring

First introduced in Ref. [30], building-scale load disaggregation is referred to as *non-intrusive load monitoring* (NILM) or sometimes *energy disaggregation*. The purpose of NILM is to monitor a building's load at the appliance level without the necessity of installing dedicated sensors on every single device—which in this case is referred to as intrusive load monitoring (ILM)—but with a unique building-level energy meter, thereby reducing costs and addressing privacy issues. While this work is not about NILM, it's actually solving a similar load disaggregation problem at a higher level of the network, with the difference that it separates consumer categories rather than individual appliances. Because of this conceptual similarity, we believe it's relevant to compare our methodological approach to established NILM techniques.

According to a recent review by Ref. [31], NILM methods fall into three primary families: machine learning, pattern matching, and source separation. It's important to note that most of these approaches, especially those in the first two families, are supervised, meaning they require access to ground truth, appliance-level measurements to which a learning procedure is applied. In contrast, our method operates in an unsupervised fashion, directly estimating sector load profiles from a composite signal without needing any ground truth measurements.

In the source separation methods, three different categories are distinguished: independent component analysis (ICA) [32–34], non-negative matrix or tensor factorization (NMF/NTF) [35–39] and sparse component analysis (SCA) [40–42]. All three methods have in common that they solve a low-rank matrix (or tensor) approximation problem, differing however in the assumptions they rely upon and the factor matrices properties they look for. A key point is that they are in practice unsuccessful when used in a fully blind, i.e., unsupervised, setting, as in this case there is often non-uniqueness of the solution and high dependence on the initialization strategy. In fact, Ref. [31] states that source separation methods developed for NILM, despite being originally thought of as unsupervised, are actually “semi-supervised” learning procedures, as they all require some degree of prior knowledge which can range from ground truth source signal data to little contextual or structural information on the sources or the mixture.

Notably, both Ref. [42] (SCA-based) and Ref. [37] (NMF-based) rely on training data consisting of individual appliance-level measurements. Their methods involve modeling each source load (appliance) using basis functions—predefined in Ref. [37] and learned in Ref. [42]—which are then used to disaggregate the aggregate load. Interestingly, Ref. [37] formulates an NMF problem with a linear constraint on the weights matrix, similar in form to ours. However, they treat this as a non-negative least squares (NNLS) problem, updating only the weights while keeping the basis fixed.

Finally, Ref. [43] proposes a contextually-supervised source separation framework. They incorporate prior information about appliances load via penalty terms, allowing for weak supervision in the absence of full training data. However, their paper does not provide details on how the resulting penalized matrix factorization problem is solved algorithmically.

2.2.2. Beyond the building level

To the best of our knowledge, only two works tackled a load disaggregation problem beyond the building scale, among which only one concerns a disaggregation in consumer categories:

- A recent study² [44] proposes an unsupervised method for separating the weather- and calendar-dependent components of the national load curve. Similarly to us, they don't rely on the observation of disaggregated measurements for model training. Nevertheless, their methodological framework does not resemble ours as they estimate the separate components as non-linear functions of weather and calendar variables with a deep-learning method. The authors themselves note a key limitation: the resulting decomposition lacks meaningfulness and practical utility. In contrast, our approach's decomposition is both identifiable and meaningful, as prior research on load profiling has shown.
- Another, more relevant work is Ref. [12], which proposes a source separation framework to estimate consumer category profiles from transformer substation load data. Their method uses a customer database containing annual consumption for each consumer category, which is available for every area covered by the substations. The main difference on the application side regards the fact that they work at the low-medium voltage vs high-voltage network scale in our case. On the methodological side, since they already know the consumer category weights in the aggregate curves, they solve a non-blind source separation problem—which boils down to NNLS—for which a global minimum exists. Our source separation problem instead is truly blind, as consumer category weights are unknown and must be estimated simultaneously with the source profiles. This optimization problem is more complex as it is non-convex and therefore no global minimum can be found [45]. Though valuable for the specific problem tackled by the authors, their study has several limitations affecting the validity of results and its generalizability to similar yet different load disaggregation problems like ours:

1. It relies on a non-typical customer database—belonging to the French DSO—providing very local consumer information which may not be available for other applications. Our method, conversely, uses widely available country-level annual and monthly consumption data.
2. As noted by the authors in the conclusion, they do not have an automatic way to determine the number of source profiles that can be extracted from the aggregate curves. In our case, though we look for a three-level decomposition, we model the aggregate load with a higher number of source profiles, directly inferred from the data.
3. By assigning a single and fixed profile to each consumer category, they assume that average consumption behaviors at this geographical scale are temporally and spatially invariant. However, Ref. [18] highlights that, while this is a reasonable assumption for households, the industry profile is very likely to differ between areas depending on the local industrial mix. Moreover, the results of Ref. [15] clearly show that household and service profiles change over the year. The framework of the present work, instead, enables flexible profiles that can account for intra-week and seasonal variations.
4. They only test the ability of their model to accurately represent a new transformer aggregate load curve but they don't validate the accuracy of their decomposition. This is obviously due to the lack of ground truth measurements, like us, but in our case we validate the disaggregated estimates at the month level against sector-level consumption indicators and compare the load profiles with those reported in the literature.

2.3. Constrained NMF

Beyond contributions specifically addressing load disaggregation problems, several *constrained* NMF (CNMF) methods have been proposed

² In which the authors confirm the literature gap we identify.

in the literature, most of which extend the original multiplicative updates (MU) algorithm introduced in Ref. [20]. A recent review can be found in Ref. [46]. However, most papers tackle highly specific cases which are not generalizable to our problem. For instance, Refs. [47,48] introduce a so-called CNMF method which is in reality more akin to nonnegative matrix *tri*-factorization (NMTF).

Several contributions—mainly in the hyperspectral unmixing literature—focus on *sparse* or *smooth* NMF (which can both be considered as belonging to CNMF, c.f. [46]), where an L^p -norm or a more elaborate regularization term is incorporated into the objective function [49–52] to promote sparsity or smoothness in the factor matrices. In a different direction, Ref. [53] propose a specific CNMF method—classified as NMF *in manifolds* in Ref. [46]—aiming at preserving pairwise similarities between observations in the reduced space. Finally, Ref. [54] provides theoretical grounding for our LCNMF method by showing that the MU algorithm can be trivially tweaked to solve NMF with additional linear constraints on the factor matrices, although the modified MU algorithms are not explicitly detailed in their paper.

Consequently, this work introduces a novel algorithmic framework for solving a general NMF problem with linear constraints on the factor matrices.

2.4. Research gap and main contributions

All aforementioned studies on aggregate load modeling and profiling rely on the availability of large, costly, end-user load samples or atypical consumer information datasets, making them poorly generalizable. Besides, most of them put efforts into the definition of granular load-based consumer classifications while failing to account for the temporal variations of daily profiles. Moreover, none of them address the problem of monitoring the total consumption of an entire consumer group, such as households, services or industry at the national level and at high frequency. Despite studies showing that highly aggregate load curve variations reflect a dynamic composition of distinctive sector consumption profiles, no work has explored the potential of load disaggregation methods for large-scale sector consumption monitoring. Finally, no existing methodological framework makes it possible to solve a general NMF problem with linear constraints, such as the one considered in this paper. To address these limitations, this paper makes the following contributions:

1. We are the first to propose a transmission-level load disaggregation method for monitoring sectoral power demand from national high-voltage grid load measurements.
2. We introduce a novel top-down approach for identifying country-level sectors daily load profiles and their temporal variations which does not require the availability of consumer-level samples and relies on open-access TSO data only.
3. We provide an algorithmic framework to solve a general NMF problem with linear constraints on the factor matrices, which can be applied to unsupervised load disaggregation problems at any scale of the power grid.

3. Methodology

Fig. 1 provides an input–output diagram of the load decomposition process.

3.1. Data

The consumption data we aim to decompose come from the Italian TSO's (Terna) hourly measurements of the national load in Italy between 2021 and 2023 [55]. Additionally, available information on sectors' electricity consumption is used. Specifically, we use two sector-level data sources:

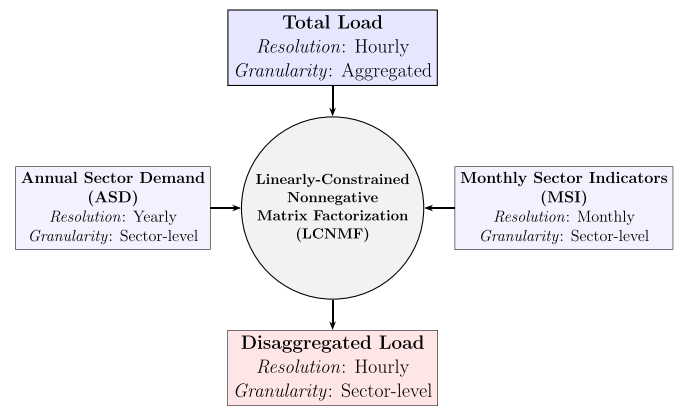


Fig. 1. Diagram showing the load decomposition process. The input data are represented in blue rectangular boxes (lighter blue for secondary inputs) while the output of the LCNMF disaggregation procedure is represented as a red rectangular box.

Table 1

Data sources and preprocessing steps.

Data	Resolution	Unit	Source	Preprocessing
Load	Hourly	MW	[55]	Daylight saving time (DST) adjustment ¹
ASD	Yearly	MWh	[56]	Merge of Agriculture into Industry and rescaling to match load data ²
MSI	Monthly	None	[57–59]	None

¹ The second occurrence of the duplicate timestamp (02:00) on hour gain days is deleted and the missing timestamp (02:00) on hour loss days is added with the load value observed at 01:00.

² Note that the decomposition, originating from load data, will inherently overestimate the sectors' consumption as it disaggregates a quantity that includes network losses, being computed from the generation side of the grid. To address this issue, one can either subtract the losses to the load data or add the losses proportionally to the ASD, before running the separation process.

- (i) The yearly consumption per sector from Terna's annual consumption report [56], which we refer to as the *annual sector demand* (ASD).
- (ii) Three unitless indicators [57–59] used by Terna, describing the monthly evolution of the household, industry³ and services sectors consumption, designated as *monthly sectoral indicators* (MSI).

Table 1 shows a tabular view of the three different data sources and the preprocessing steps applied.

3.2. Qualitative description

Blind source separation (BSS) is a technique used to separate a set of mixed signals into their original, independent source components, without prior knowledge of the mixing process or the source signals. In our load disaggregation problem, we observe each day a load curve having a specific shape resulting from the mixture of the sectors' load profiles—which can be seen as the source signals. We aim to leverage the fact that some sectors are predominant on certain days to identify their load

³ This also includes the agricultural sector. Given that the agricultural sector's weight in total consumption is negligible compared to the three other sectors, and since we have no suitable indicator to act as a proxy for its monthly consumption, we have chosen to merge this sector with the industrial one (as we believe their consumption patterns are most similar).



Fig. 2. Illustration of the load disaggregation problem. In this example, each sector is associated with one source only.

profile, estimate their contribution to the overall mixture, and reconstruct their load signal. However, beyond this necessity to have mixture weights, i.e., sector importance, that show some degree of variability, BSS assumes that sectors have distinct and time-invariant load profiles. While it is reasonable to expect that the typical load curve of one sector differs from another, it is clear that these profiles can vary based on factors such as the day of the week (weekday vs. weekend) or the season (winter vs. summer) as shown in Ref. [15]. Therefore we cannot consider the sectors' load profiles as fixed and constant over time. To address this limitation, we assume that each sector's load profile is itself a mixture of multiple underlying sources, which may or may not be directly interpretable, and that these sources meet the conditions for separation. Note that the number of sources present in the aggregated signal can be inferred from the data with dimensionality reduction techniques such as principal component analysis (PCA). By identifying the contribution of these sources to the overall load signal, we can estimate the contribution of each sector by aggregating the corresponding sources (Fig. 2).

However, without incorporating prior information on the sectors profiles or weights, the likelihood of finding the desired decomposition—where the sources are precisely the “voices” of our sectors—is quite low. Fortunately, we do have valuable data that can guide the separation process: the ASD and MSI. The ASD provides information on sectors' total annual consumption, while the MSI indicates how this consumption is distributed over the 12 months. Hence an estimate of each sector's monthly consumption can be computed, and this can be compared to the results of the decomposition, aggregated at the month level. By constraining the decomposition to produce monthly consumption estimates that align with these known values, we can identify sources that are hierarchically linked to the sectors and hence reconstruct the sectors' load.

Once we have identified the sources, any *new* set of daily load curves can be decomposed into contributions from these sources without requiring any further information about sectors' consumption (which, by contrast, is necessary for the initial source identification). In other words, we can decompose new load data using only the load curves themselves (and the identified sources, obviously). The source separation process becomes “non-blind” because we only need to estimate the weights of the observed mixtures. To validate this nowcasting predictive approach, we decompose the load curves for the year 2023 using the sources estimated from 2021-2022, without using any further sector-level data. We then compute the predicted monthly consumption and compare it to the MSI for verification.

The next two subsections delve into the mathematical details of what was just described.

3.3. Model

Let $L_i : [0, T] \rightarrow \mathbb{R}^+$ represent the system's⁴ load as a function of time during a day $i \in \{1, \dots, n\}$. The total energy consumption e_i during day i is given by:

$$e_i = \int_0^T L_i(t) dt = \|L_i\|_1$$

We define $X_i : [0, T] \rightarrow \mathbb{R}^+$ as the load shape curve observed on day i :

$$L_i = e_i \cdot X_i$$

In other words, X_i is the normalised load curve of day i such that it corresponds to a total consumption of 1. For notational convenience, we may assimilate X_i to an element of \mathbb{R}_+^p , corresponding to evaluations of the underlying function at p regularly-spaced time instants spanning the interval $[0, T]$. For example, $p = 25$ corresponds to evaluations at 00:00, 01:00, ..., 24:00.

Each curve X_i is assumed to be a mixture of K sources S_1, \dots, S_K with $K \ll n$:

$$\begin{cases} X_1 = c_{11}S_1 + c_{12}S_2 + \dots + c_{1K}S_K + \epsilon_1 \\ X_2 = c_{21}S_1 + c_{22}S_2 + \dots + c_{2K}S_K + \epsilon_2 \\ \vdots \\ X_n = c_{n1}S_1 + c_{n2}S_2 + \dots + c_{nK}S_K + \epsilon_n \end{cases} \quad (1)$$

where for $i \in \{1, \dots, n\}$ and $k \in \{1, \dots, K\}$, the coefficient c_{ik} represents the *concentration* of source k in the load (shape) curve X_i and $\epsilon_1, \dots, \epsilon_n$ are additional noise terms. The profiles S_1, \dots, S_K and the concentrations c_{ik} 's are constrained to be positive. Additionally, we impose that S_1, \dots, S_K integrate to 1 (as the X_i 's) and thus $c_{i1} + \dots + c_{iK} = 1$, for all $i \in \{1, \dots, n\}$. The system of Eq. (1) can be re-written in a matrix form:

$$X \approx CS \quad (2)$$

where X is the $n \times p$ load curves matrix, C the $n \times K$ concentrations matrix and S the $K \times p$ sources matrix:

$$X = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix}, C = \begin{bmatrix} c_{11} & c_{12} & \dots & c_{1K} \\ c_{21} & c_{22} & \dots & c_{2K} \\ \vdots & \vdots & \dots & \vdots \\ c_{n1} & c_{n2} & \dots & c_{nK} \end{bmatrix} \text{ and } S = \begin{bmatrix} S_1 \\ S_2 \\ \vdots \\ S_K \end{bmatrix}.$$

The load can hence be decomposed by finding a pair of matrices \hat{C} and \hat{S} whose product $\hat{C}\hat{S}$ best approximates the data X . The estimated

⁴ The national power grid in our case

consumption \hat{e}_{ik} of source k during day i is then deduced by multiplying the estimated concentration \hat{c}_{ik} with the total consumption e_i :

$$\hat{e}_{ik} = \hat{c}_{ik} e_i$$

K can be chosen by estimating the underlying dimension d of the space in which the load curves X are embedded, e.g. with functional principal component analysis (FPCA) [60]. Given that the fitted curves CS are convex linear combinations of the sources and therefore live in a $(K - 1)$ -dimensional linear subspace of \mathbb{R}_+^p , we must choose $K = d + 1$.

The low-rank matrix approximation problem of equation (2), when both matrices are constrained to have all their coefficients positive, is known as *Non-negative Matrix Factorization* (NMF), with the difference that we have an additional constraint due to the fact that the profiles contained in S must integrate to 1. Moreover, as stated in the previous section, we would like to add a constraint on the solution such that it is consistent with the ASD-MSI sectors' monthly consumption estimates. To formalise this, let $E = [e_1 \ e_2 \ \dots \ e_n]'$ be the n -dimensional vector of daily consumption, for which we consider a block-form notation in which the vector elements are grouped by month:

$$E = \begin{bmatrix} E_1 \\ E_2 \\ \vdots \\ E_m \end{bmatrix}$$

where m is the number of months in the period considered. We additionally define a $m \times n$ matrix B and re-write the $n \times K$ matrix C of concentrations in a similar block-form manner:

$$B = \begin{bmatrix} E'_1 & 0 & \dots & 0 \\ 0 & E'_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & E'_m \end{bmatrix} \text{ and } C = \begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_m \end{bmatrix}$$

Multiplying the two matrices gives us the $m \times K$ matrix of the monthly consumption of the K sources:

$$BC = \begin{bmatrix} E'_1 C_1 \\ E'_2 C_2 \\ \vdots \\ E'_m C_m \end{bmatrix}$$

The problem now concerns mapping the K sources to the g sectors. We may assign more sources to sectors where we expect the characteristic load profile to vary more significantly throughout the year. We formalise this mapping by defining a surjection $\sigma : \{1, \dots, K\} \rightarrow \{1, \dots, g\}$ and the associated $K \times g$ matrix A such that, for source $i \in \{1, \dots, K\}$ and sector $j \in \{1, \dots, g\}$:

$$A_{i,j} = \begin{cases} 1 & \text{if } j = \sigma(i) \text{ i.e., if } i \text{ is a source of sector } j \\ 0 & \text{otherwise} \end{cases}$$

For instance, in the case $K = 5$ and $g = 3$, attributing source 1 to sector 1, sources 2 and 3 to sector 2, and sources 4 and 5 to sector 3, we have:

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}$$

We can finally obtain the $m \times g$ matrix Z of the estimated monthly consumption per sector:

$$Z = BCA$$

On the other side, let $Y = [Y_1 \ \dots \ Y_g]$ be the $m \times g$ matrix whose columns Y_j 's represent the MSI, which, as highlighted in Section 3.2, only carry

information on the *shape* that the columns of Z , seen as monthly time series, must assume. Intuitively, we would like to drive the decomposition to yield a Z matrix with columns as *correlated* as possible with those of Y , regardless of the difference in scale. However, the introduction of a penalty term based on the Pearson correlation coefficient—or other relevant similarity measure quantifying statistical dependence—between Z_j and Y_j in the loss function may be computationally intractable. To address this issue, our approach is to use the ASD to *rescale* the matrix Y so that it can directly be compared to Z . More specifically, let w_j , for $j \in \{1, \dots, g\}$, be the total consumption given by the ASD of sector j over the period considered (which must necessarily be composed of entire years). Note that we should have $w_1 + \dots + w_g = \|E\|_1$. Performing the following scaling operation, for all $j \in \{1, \dots, g\}$:

$$Y_j \leftarrow \frac{Y_j}{\|Y_j\|_1} w_j$$

Y is now exactly an estimate of what Z should be. Note that the boundary condition $Y_1 + \dots + Y_g = [\|E_1\|_1, \|E_2\|_1, \dots, \|E_m\|_1]'$, i.e., summing the sectors' consumption should give the total consumption at the month level, may not be fully satisfied but we can similarly rescale Y row-wise to correct this effect.

We can write down the final optimisation problem:

$$\begin{aligned} (\hat{C}, \hat{S}) = \arg \min_{C \in \mathbb{R}_+^{m \times K}, S \in \mathbb{R}_+^{K \times p}} \text{dist}(X, CS)^2 \\ \text{s.t. } BCA = Y \\ S \mathbf{1}_p = \mathbf{1}_m \end{aligned} \quad (3)$$

where $\text{dist}()$ is some distance measure on $\mathbb{R}_+^{n \times p}$ and $\mathbf{1}_r$ is the r -dimensional vector of ones.

3.4. Solution method

The Multiplicative Updates (MU) algorithm introduced by Lee and Seung [20] allows one to find a local minimum to the ordinary NMF problem, formulated with a distance measure being the metric induced by the Frobenius norm (or another distance measure similar to the Kullback-Leibler divergence). Essentially, starting from arbitrary initial values, the matrices C and S are alternately updated by multiplying their current values by a factor that depends on the quality of the approximation, until convergence.

We propose a modified version of the MU algorithm to solve the *linearly constrained* non-negative matrix factorization (LCNMF) problem, i.e., NMF with additional equality constraints that take a linear form in the factor matrices, generalizing (3). Indeed, given any $B \in \mathbb{R}_+^{m \times n}$, $A \in \mathbb{R}_+^{K \times g}$, $Y \in \mathbb{R}_+^{m \times g}$, $F \in \mathbb{R}_+^{l \times K}$, $D \in \mathbb{R}_+^{p \times q}$, $Z \in \mathbb{R}_+^{l \times q}$, let

$$\begin{aligned} (\hat{C}, \hat{S}) = \arg \min_{C \in \mathbb{R}_+^{m \times K}, S \in \mathbb{R}_+^{K \times p}} \text{dist}(X, CS)^2 \\ \text{s.t. } BCA = Y \\ FSD = Z. \end{aligned} \quad (4)$$

We reformulate (4) by choosing the metric induced by the Frobenius norm as the distance measure and translating the constraints in terms of distance. Considering then its Lagrangian form, we obtain the following unconstrained loss function to be minimized, with constraints incorporated as penalization terms:

$$\mathcal{L}(C, S) = \|X - CS\|_F^2 + \alpha \|BCA - Y\|_F^2 + \beta \|FSD - Z\|_F^2 \quad (5)$$

for some $\alpha, \beta > 0$ penalizing violations of the constraints.

To minimise \mathcal{L} , we start from arbitrary initial values and update the factor matrices as follows, until convergence:

$$S \leftarrow S \cdot \frac{C'X + \beta F'ZD'}{C'CS + \beta F'FSDD'}, \quad C \leftarrow C \cdot \frac{XS' + \alpha B'YA'}{C'SS' + \alpha B'BCAA'}$$

where \cdot (multiplication) and — (division) operations are coefficient-wise. These updates are derived following the same logic as in Ref. [20], further detailed in the appendix.

A first consideration regards the determination of suitable α and β values as they define the weight that will be given to the terms penalizing the violation of the constraints. Given that the penalization loss terms may not be of the same magnitude as the factorization loss term, their primary purpose is a rescaling one. Analysing jointly the three loss terms at convergence obtained for different values of (α, β) , one can determine suitable values.

A second consideration concerns the initialization strategy, which is obviously a critical point. Our approach involves running the MU algorithm $N = 1000$ times, each time with a different initialization—keeping the same initial value $\left[\frac{1}{p} \dots \frac{1}{p}\right]'$ for the rows of S but obtaining each time a different initial value for C by uniformly sampling its rows on the simplex $\mathcal{S}(K) = \{x \in \mathbb{R}_+^K \mid x_1 + \dots + x_K = 1\}$ —generating a large set of solutions. This set of solutions is then clustered based on the loss function value at convergence and the cluster of solutions associated with the minimum value is retained. This allows us to work with an ensemble of N^* solutions that can be combined to obtain more robust estimates and to quantify uncertainty. The ensemble estimated consumption \hat{e}_{ik} of source k during day i is hence deduced by averaging the individual estimates given by the N^* solutions:

$$\hat{e}_{ik} = \frac{1}{N^*} \sum_{l=1}^{N^*} \hat{c}_{ik}^{(l)} e_i$$

Additionally one can also retrieve source k 's estimated load during day i :

$$\hat{L}_{ik} = \frac{1}{N^*} \sum_{l=1}^{N^*} \hat{c}_{ik}^{(l)} e_i S_k^{(l)}$$

As mentioned in Section 3.2, we can project a new load curve X_0 on the linear space spanned by the solution matrix \hat{S} to estimate the sources concentrations hence sectors' contributions to X_0 . The source separation problem becomes “non-blind” as we only need to estimate the concentrations:

$$\hat{C}_0 = \arg \min_{C \in \mathbb{R}_+^{n \times K}} \|X_0 - C\hat{S}\|_F^2$$

This is a convex optimisation problem for which the global minimum can be found, as long as the sources are linearly independent (i.e., \hat{S} is full rank). In practice, we use the standard MU of Ref. [20] to solve the problem, updating C only.

4. Results

4.1. LCNMF solutions analysis and interpretation

We applied the source separation algorithm to 2021 and 2022 data, acting as a training set, to identify the sources and consequently the sectors' profiles. We held 2023 data out as a test set on which we estimated the concentrations of the sources identified in the training set, without incorporating the ASD-MSI information. We then compared the sectors' monthly consumption of 2023 estimated from the decomposition with the ASD-MSI target.

Fig. 4 shows the scree plot of the FPCA performed on 2021–2022 load curves. We can see that 4 principal components explain $\approx 97\%$ of the variance and appear to be a good choice for dimensionality reduction. We can therefore deduce that we should achieve a satisfying decomposition with $K = 5$ sources. Trials clearly indicate that the optimal sources/sectors mapping is obtained by assigning two sources to households, two sources to services and one source to industry.

Concerning the penalization parameters, we chose $\alpha = 3 \cdot 10^{-10}$ and $\beta = 1$ as they provided the best trade-off between matrix factorization quality and compliance with the constraints. Following the $N = 1000$




Fig. 3. Value of the loss \mathcal{L} at convergence for the $N = 1000$ LCNMF runs. The values are vertically jittered for legibility. The grey dashed line indicates the threshold used to isolate the red cluster.




Fig. 4. Scree plot of the FPCA run on 2021–2022 load curves.

Table 2

Comparison of the low-rank matrix approximation quality between FPCA and LCNMF.

Loss	$\ A - B\ _1$	$\ A - B\ _F$	$\ A - B\ _\infty$
FPCA	9.2	0.0096	0.0053
LCNMF	8.9	0.0077	0.0051

LCNMF runs, two compact and well-separated clusters of loss values at convergence were observed. We then simply used a threshold of $\mathcal{L} = 0.01$ to retrieve the cluster of $N^* = 271$ optimal solutions (Fig. 3).

As underlined in Section 3.4, we do not choose a specific solution but we analyse and use the entire set of N^* optimal solutions to obtain a robust ensemble estimation and quantify uncertainty.

A first consideration regards the quality of the approximation of the load curves that we obtain with the LCNMF decomposition. Table 2 compares the average low-rank matrix approximation quality of LCNMF with the approximation provided by the projection of the curves onto the space spanned by the first 4 functional principal components.

The comparison is relevant because both approximation methods rely on a 4-dimensional representation. We observe that LCNMF provides a better fit than FPCA, showing that the approximation quality is more than satisfactory.

The extracted sources are shown in Fig. 5. We notice considerable variability across the different solutions, except for the source corresponding to the industrial sector. In this case, all solutions converge on the same load profile, marked by a sharp morning ramp beginning at 05:00 and three distinct peaks at 09:00, 16:00, and 20:00. Sources mapped to the household sector show a morning ramp that either starts steeply at 05:00 or more gradually at 08:00. The first one displays an evening peak around 18:00–19:00, while the second one shows it around

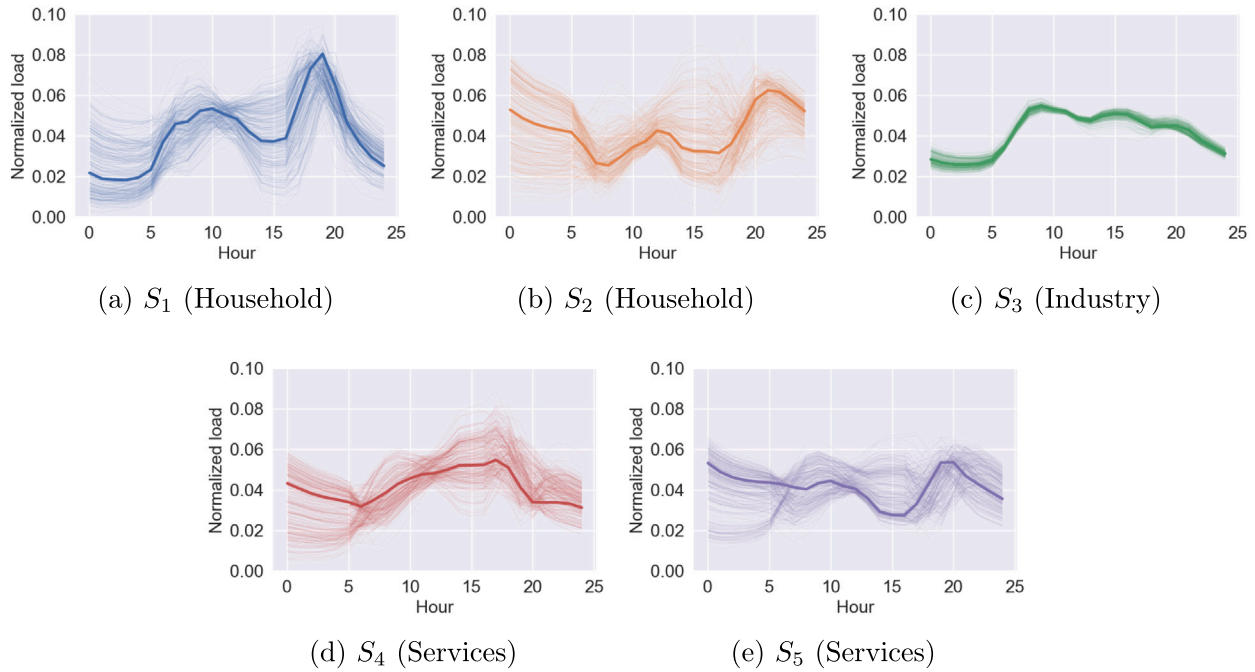


Fig. 5. Sources corresponding to the $N^* = 271$ retained solutions. Each thin line corresponds to one solution while the thick line, added for visualization purposes, is the geometric medoid of the thin lines, which is the closest observation to the geometric median (in terms of Euclidean distance). Since inside a sector the source number is arbitrary, the sources of household and services were aligned on the order given by the first solution.

21:00–22:00. Beyond these few observations, the results are hard to interpret. The sources linked to the services sector are generally flatter than those of the household sector. They typically exhibit a morning ramp either at 05:00 or 07:00, depending on the solution. In particular, for most solutions, S_4 (Services) shows an afternoon peak around 17:00–18:00. S_5 (Services) seems to be characterized, in most cases, by a rather constant regime with an afternoon trough centered on 15:00 and an evening peak located around 20:00.

Overall, except for the industrial sector whose load profile is directly described by S_3 (Industry), the interpretability of the resulting sources for the two other sectors is limited as their load profile is given by a positive linear combination of their sources. For these two sectors, there is in fact a profile specific to each day as the weights of the linear combination are free to vary, thereby causing a modification of the profile whenever one weight increases in proportion to the other. Fig. 6 shows the average profile of households and services that we obtain depending on the type of day—Monday, Working day, Saturday and Holiday⁵—and the season.⁶ It is quite clear that enabling flexible profiles allows us to model weekly and yearly seasonal variations in these two sectors’ daily consumption patterns. Specifically, we can see that the households profile is represented by a double-hump curve with a first small hump at 10:00–11:00 and a second higher one around 19:00–20:00. The curve is flatter and has a shallow and late morning ramp during the weekends compared to the weekdays. We can also note that the evening peak is delayed by one hour, occurring more around 20:00 than 19:00. Concerning the seasonal effect, we observe a fall-winter vs a spring-summer regime. Fall-winter profiles show higher amplitude, an earlier and steeper morning ramp, and an earlier evening peak around 19:00. On average, the services profiles show a nearly constant activity between 10:00 and 20:00 with no clear morning or evening peak, which is consistent with business

and office hours in Italy. While the day type effect is not so clear, we observe a seasonal effect that is similar to that for households, with the difference that summer and spring show slightly diverging behaviors in the afternoon.

The interpretation of the identified sources is made easier by looking at their concentrations measured every day. Specifically, we can try to understand whether a source is more or less active depending on the day type and the period of the year.

Fig. 7 shows the distribution of the average concentration of each source per type of day and season. We can therefore appraise both the effect of day type/season, and the variability, or disagreement, among the solutions. Overall, the day type seems to have a clear effect on the activity of any source (Fig. 7 a): on average, all five sources contribute equally to holidays while Mondays and working days are clearly dominated by S_3 (Industry). Saturdays lie in-between. Besides, we can note the strong agreement between solutions—which can be called “models” at this point—on the average concentration estimated during working days compared to the other day types. On the contrary we observe considerable disagreement in the disaggregation of Mondays’ load curves, especially regarding S_3 ’s contribution.

For what regards the seasonal effect, Fig. 7(b) highlights a clear opposition between spring-summer and fall-winter conditions for sources associated with households, S_1 and S_2 , as was observed previously in the profiles. This observation supports the existence of two load regimes characterizing household consumption, one under cold temperatures, driven by S_1 , and one under warm temperatures, driven by S_2 . A possible explanation is the quadratic effect of temperature on power consumption in Italy [17] with electric heating in winter and air conditioning devices in summer, causing two distinct seasonal shifts in the daily consumption profile. Note that this dual effect may not apply to countries where there are only either heating or cooling needs, e.g., Denmark. The seasonal effect is also pronounced on S_3 (Industry), with the lower activity in summer (due to factory closures in August) and a highest activity in spring, indicating therefore that this effect is more a calendar than a climatic one. A summer effect opposed to the spring one is also perceived for S_4 (Services).

⁵ Sundays and Bank holidays

⁶ The mapping is made at the month level so winter is composed of all days from January to March, spring from April to June, summer from July to September and fall from October to December.

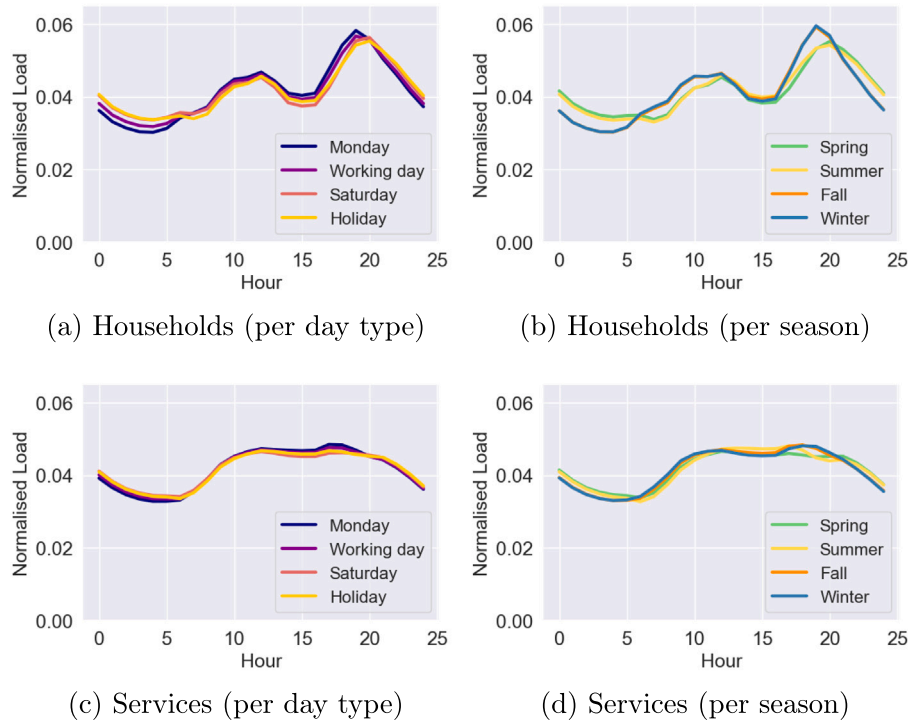


Fig. 6. Average load profiles found for the households and services sectors depending on the day type and the season. The load profiles were first averaged at the day type and season level for each solution separately after which the average across solutions was retrieved. For legibility, we do not show the result for all individual solutions.

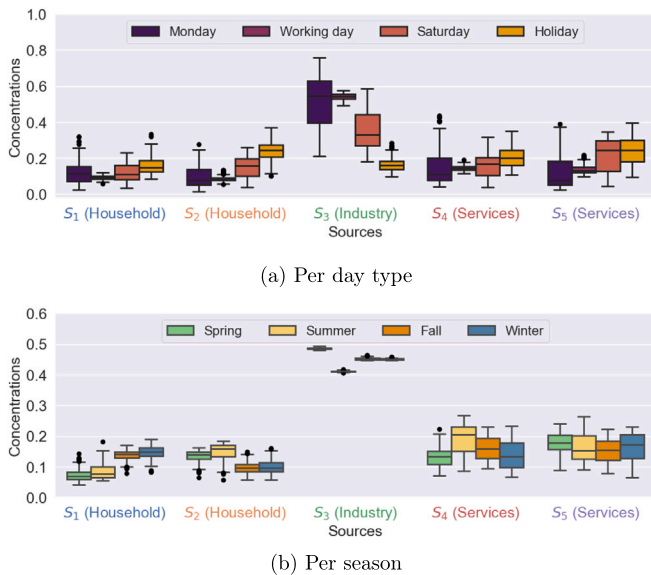


Fig. 7. Distribution of the average concentrations per day type and season across solutions.

4.2. Disaggregated consumption

The core output of the process, as indicated in Fig. 1, is the disaggregated load signal, with the three separate contributions of the sectors. We first compute the daily load curve of each source, obtained from the shape curve rescaled with the fraction of the daily consumption corresponding to that source, itself obtained by multiplying the concentration with the daily consumption. We then add up the sources curves to get the sectors curves, which we finally concatenate along days to get a

continuous signal. Note that the signal presents an inherent discontinuity every day at midnight since the concentration varies non-smoothly from day to day. This discontinuity is obviously a mathematical artifact, which however does not affect the aggregated daily quantities. Nonetheless, for applications where this discontinuity is deemed to be problematic, it can be overcome by smoothing the concentrations regarded as hourly data over time, keeping daily average concentrations equal to the original ones. In this paper, we don't pursue the matter further.

The disaggregated load signal for the first week of February 2022 is shown in Fig. 8(b). Beyond the uncertainty which is quite high for the household and services sectors compared to the industrial one, we can observe that the uncertainty is higher during the weekend and even higher on Monday. This confirms the earlier observed fact that Mondays have a different status from the other weekdays.

Unfortunately, we don't have sector-level consumption data to validate the results of the disaggregation at the hourly, daily or weekly level. But moving to the monthly level, we have the MSI to which our estimates can be compared, which are the data, together with the ASD, that were used to guide the separation process. As 2023 was left out in the blind separation process—meaning that the sector-level information corresponding to that year is completely unknown to the model—we can reliably validate our disaggregation procedure by verifying whether monthly sectors' consumption estimates align with the MSI for that year. As a benchmark, we include a naive forecast corresponding to one-year-lagged values and a yearly-seasonal exponential smoothing model: Holt-Winter's (HW) additive method with no trend⁷ [61]. For the HW model, we forecast the 12 months of 2023 using 2021–2022 data only—similarly to the BSS approach.

The comparison of these out-of-sample predictions with the MSI for 2023 is reported in Fig. 9 and Table 3. All three estimation methods

⁷ Given the series are stationary.

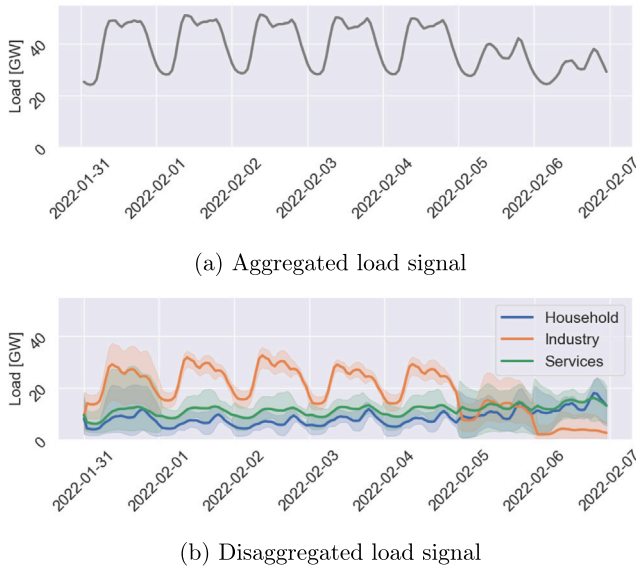


Fig. 8. Result of the load disaggregation for the week starting on Monday January 31st, 2022. Shaded areas on the bottom panel represent the one-at-a-time 95% prediction intervals given by the 0.025 and 0.975 quantiles of the ensemble estimation.

align well with the observed values. While BSS appears to reproduce the dynamics of the observed series slightly better than the naive forecast, its performance remains comparable to that of HW. These results suggest that the BSS approach produces accurate monthly consumption estimates and is competitive for nowcasting sectors monthly consumption.

5. Discussion

Though designed for the specific problem tackled in this work, the proposed LCNMF method is highly generalizable as it allows solving a load disaggregation problem at any scale of the power network, for which prior information—such as low-frequency disaggregated consumption values or direct information on one or more load

Table 3
Pearson’s correlation coefficient with MSI for the test year 2023.

	Naive forecast	HW forecast	BSS
Household	0.914	0.959	0.951
Industry	0.962	0.961	0.974
Services	0.912	0.971	0.970

components—can be expressed as linear constraints on the factor matrices. While the current formulation addresses linear constraints only, the framework can be easily adapted to support more complex non-linear forms, provided they involve functions differentiable with respect to the factor matrices (see Appendix A).

The application of our BSS framework to Italian national load data revealed established features of household, industry, and services sectors’ electricity consumption profiles. There is in fact a strong consistency between the profiles we estimate and those identified in previous studies relying on ground truth data [3,12,15], thereby supporting the validity of our results and the effectiveness of the proposed method. In particular, households consumption shows a daily profile that is dominated by an evening peak and is driven by two distinct regimes, in warmer vs colder seasons. The services sector is found to be especially active during the usual business and office hours and shows a singular summer daily load profile, which could be related to the air conditioning devices working at their maximum around 16:00, and only in summer. These patterns globally align with those observed in Ref. [15] for Denmark, with some nuances likely due to climate and lifestyle differences. Finally, industrial consumption is particularly low during holiday periods and does not seem to be much affected by seasons. The trimodal pattern we identify for the industrial daily load profile is very similar to that reported by Ref. [12,15]. Besides, the fact that a single source satisfactorily describes industry’s daily load profile, unlike household and services profiles which are described by linear combinations of two sources each, is in line with the findings of Ref. [15] on the invariance of industrial power consumption. This consistent daily consumption pattern allows for a clear distinction of industry from the other two sectors. In contrast, households and services profiles appear to show a certain degree of similarity that makes them harder to dissociate. Besides, the monthly sectoral consumption estimates derived from disaggregated load were found to closely align with the TSO’s official statistics—indices based on consumer samples—and be a competitive approach for nowcasting sectors’ monthly consumption.

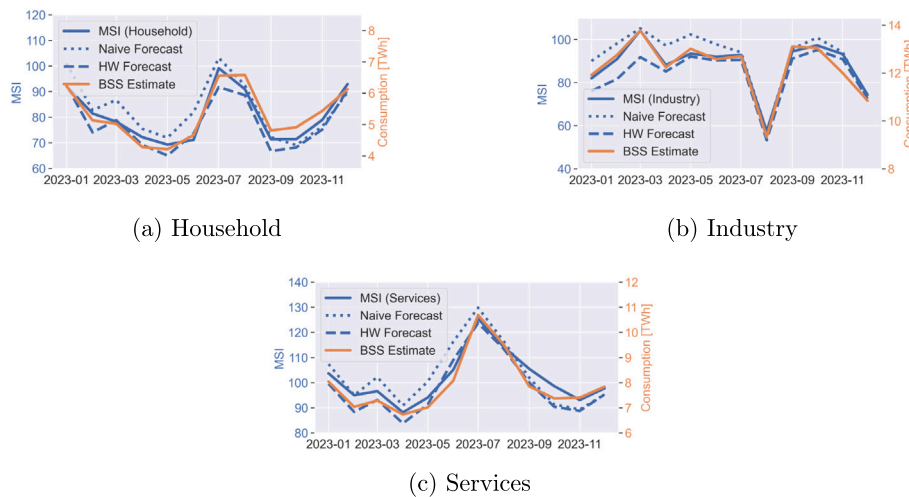


Fig. 9. Comparison of the BSS monthly sector consumption estimates with the TSO’s MSI and its naive and HW forecasts for the test year 2023. Note that the data are not on the same scale and are therefore represented with two different ordinate axes (left/blue for MSI and right/orange for BSS).

A primary limitation of this application is the lack of ground truth data at fine temporal resolutions, such as hourly or daily, which would enable (i) improved supervision of the source separation process and (ii) a direct assessment of disaggregation accuracy, thereby complementing the literature-based validation. Besides, the model assumes a hierarchical relationship between the identified sources and the sectors, an assumption that may not hold in practice. Indeed, a given source may be associated with multiple sectors—for example, air conditioning usage in both residential and commercial consumption. Finally, we do not explicitly control for weather effects which are likely to introduce additional variability in the load signal. Properly accounting for such effects, by incorporating weather data directly into the model, may yield more consistent load profiles and lower estimation uncertainty. Another possible solution is to adopt a methodology to separate the weather-dependent load component, such as that of Ref. [44], prior to the sectoral separation. Future research could extend the empirical validation to other countries or smaller power zones where high-frequency ground truth consumption data are available, potentially considering a more granular consumer segmentation.

From a broader perspective, models capable of detecting and profiling consumption patterns, distinguishing among different categories of customers—such as that proposed in the present work, are a valuable tool for flexibility management. Indeed, while generation is becoming more diffuse and less predictable, power demand is evolving due to new drivers such as electric vehicles, heat pumps and data centers. Concurrently, business opportunities emerge, resulting in a much more complex energy market structure. Within this framework, system operators may lack the resources to monitor all the assets necessary to effectively manage the power system. Consequently, they could benefit from tools able to reconstruct sector-aware power profiles with appropriate granularity. This possibly unlocks opportunities to efficiently schedule balancing actions, such as shifting, curtailment and modulation of loads, thereby reducing peak generation and the need for costly reserves. Moreover, the use of advanced, standardized profiling models for end-user demand can enhance TSO-DSO coordination, improving system-wide balancing and power system resilience.

Another central takeaway from this work is the potential of statistical decomposition methods for solving electrical load profiling or disaggregation problems *beyond the building scale*, even when no disaggregated consumption measurements are available at high-frequency. We hope this contribution will encourage further exploration of this topic.

6. Conclusion

Monitoring electricity consumption at the level of residential, commercial and industrial sectors is of considerable interest to the TSO and other power system stakeholders. To date, no solution exists to perform it live, at country scale, high frequency and utilizing open-access national grid load measurements only. In this context, we introduced a novel signal separation framework to monitor sectors consumption from open-access aggregate grid-scale load measurements collected by the TSO. The proposed framework is highly generalizable as it allows solving a load disaggregation problem at any scale of the power network where disaggregated ground truth data are not available, provided there is some prior information on load components which can be expressed as linear constraints on the factor matrices. Applied to Italian national load data over the 2021–2023 period, our method identifies residential, commercial and industrial daily load profiles which are consistent with existing literature. Besides, sectoral monthly consumption estimates obtained from load disaggregation align with official sector-level power demand statistics used by the TSO.

The main limitation of this study lies in the absence of quantitative validation at the hourly or daily level, due to the lack of ground truth data at these levels. The application to other countries or smaller zones where granular disaggregated load data are available is left for future research.

CRedit authorship contribution statement

Guillaume Koechlin: Writing—original draft, Visualization, Software, Methodology, Conceptualization. **Filippo Bovera:** Writing—review & editing, Methodology, Conceptualization. **Elena Degli Innocenti:** Writing—review & editing, Validation, Resources, Methodology, Conceptualization. **Barbara Santini:** Writing—review & editing, Validation, Resources, Conceptualization. **Alessandro Venturi:** Writing—review & editing, Validation, Resources, Conceptualization. **Simona Vazio:** Writing—review & editing, Validation, Resources, Conceptualization. **Piercesare Secchi:** Writing—review & editing, Methodology, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT to obtain suggestions for improving writing fluency and clarity. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Funding

The present research has been partially supported by MUR, grant Dipartimento di Eccellenza 2023–2027.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors wish to thank Valeria Amoretti and Luca Marchisio of Terna S.p.A. for their contributions to the interesting and deep discussions on the topics developed in this paper.

Appendix A. Modified multiplicative updates for the resolution of LCNMF

Given $B \in \mathbb{R}_+^{m \times n}$, $A \in \mathbb{R}_+^{K \times g}$, $Y \in \mathbb{R}_+^{m \times g}$, $E \in \mathbb{R}_+^{l \times K}$, $D \in \mathbb{R}_+^{p \times q}$, $Z \in \mathbb{R}_+^{l \times q}$, we propose to solve:

$$(\hat{C}, \hat{S}) = \arg \min_{C \in \mathbb{R}_+^{n \times K}, S \in \mathbb{R}_+^{K \times p}} \mathcal{L}(C, S, \alpha, \beta) \quad (\text{A.6})$$

where

$$\mathcal{L}(C, S, \alpha, \beta) = \|X - CS\|_F^2 + \alpha \|BCA - Y\|_F^2 + \beta \|ESD - Z\|_F^2$$

and α, β are two non-negative penalization parameters.

The case $\alpha = \beta = 0$ corresponds to the ordinary NMF problem:

$$\mathcal{L}(C, S) = \mathcal{L}(C, S, 0, 0) = \|X - CS\|_F^2$$

Using matrix differential calculus, with $'$ denoting the transpose:

$$\nabla_S \mathcal{L}(C, S) = 2C'CS - 2C'X \quad \text{and} \quad \nabla_C \mathcal{L}(C, S) = 2CSS' - 2XS'$$

Separating positive and negative contributions, the gradient $\nabla_S \mathcal{L}$ or $\nabla_C \mathcal{L}$ can be written:

$$\nabla_S \mathcal{L}(C, S) = \nabla_S^+ \mathcal{L}(C, S) - \nabla_S^- \mathcal{L}(C, S)$$

The popular multiplicative updates (MU) algorithm of Lee and Seung [20], in the case of Frobenius loss, proposes to alternately update C and S , as follows:

$$S \leftarrow S \cdot \frac{C'X}{C'SS}, \quad C \leftarrow C \cdot \frac{XS'}{CSS'}$$

where \cdot and $\frac{\cdot}{\cdot}$ operations are coefficient-wise.

There are several ways to derive these updates, mentioned in the original paper by Lee and Seung and detailed in Section 8.2 of Ref. [45], among which the “gradient ratio heuristic” shows that these updates can be written:

$$C \leftarrow S \cdot \frac{\nabla_S^- \mathcal{L}(C, S)}{\nabla_S^+ \mathcal{L}(C, S)}, \quad C \leftarrow C \cdot \frac{\nabla_C^- \mathcal{L}(C, S)}{\nabla_C^+ \mathcal{L}(C, S)}$$

We can apply the same idea in the case $\alpha \neq 0$ or $\beta \neq 0$. In this case we have:

$$\nabla_C \mathcal{L}(C, S, \alpha, \beta) = 2CS'S' + 2\alpha B'BCAA' - (2XS' + 2\alpha B'YA')$$

$$\nabla_S \mathcal{L}(C, S, \alpha, \beta) = 2C'CS + 2\beta E'ESDD' - (2C'X + 2\beta E'ZD')$$

and we can therefore consider the following updates:

$$S \leftarrow S \cdot \frac{C'X + \beta E'ZD'}{C'CS + \beta E'ESDD'}, \quad C \leftarrow C \cdot \frac{XS' + \alpha B'YA'}{C'SS' + \alpha B'BCAA'}$$

With a well-chosen auxiliary function, the proofs in Ref. [20] can be directly adapted to prove that the previous updates lead to a monotonically decreasing loss \mathcal{L} and a converging update sequence.

Code and data availability

The analyses presented in this work were conducted using Python 3.11 and R 4.3.0. Code and data are available at https://github.com/gkchln/bss_load. Note that the monthly sector indicator [57] cannot be shared due to privacy reasons and is therefore not included in the repository.

References

- [1] IEA, Key World Energy Statistics 2021, Technical Report, IEA, Paris, 2021. <https://www.iea.org/reports/key-world-energy-statistics-2021/final-consumption>.
- [2] G. Mor, J. Cipriano, G. Martirano, F. Pignatelli, C. Lodi, F. Lazzari, B. Grillone, D. Chemisana, A data-driven method for unsupervised electricity consumption characterisation at the district level and beyond, *Energy Rep.* 7 (2021) 5667–5684, <https://doi.org/10.1016/j.energy.2021.08.195>, <https://www.sciencedirect.com/science/article/pii/S2352484721007988>.
- [3] N. Voulis, M. Warnier, F.M.T. Brazier, Understanding spatio-temporal electricity demand at different urban scales: a data-driven approach, *Appl. Energy* 230 (2018) 1157–1171, <https://doi.org/10.1016/j.apenergy.2018.08.121>, <https://www.sciencedirect.com/science/article/pii/S0306261918312959>.
- [4] J. Langevin, J.L. Reyna, S. Ebrahimigharebaghi, N. Sandberg, P. Fennell, C. Nägeli, J. Laverge, M. Delghust, E. Mata, M. Van Hove, J. Webster, F. Federico, M. Jakob, C. Camarasa, Developing a common approach for classifying building stock energy models, *Renew. Sustain. Energy Rev.* 133 (2020) 110276, <https://doi.org/10.1016/j.rser.2020.110276>, <https://www.sciencedirect.com/science/article/pii/S1364032120305645>.
- [5] Y. Wang, Q. Chen, C. Kang, M. Zhang, K. Wang, Y. Zhao, Load profiling and its application to demand response: a review, *Tsinghua Sci. Technol.* 20 (2015) 117–129, <https://doi.org/10.1109/TST.2015.7085625>, <https://ieeexplore.ieee.org/document/7085625>.
- [6] F. Møller Andersen, P.A. Gunkel, C. Bjerregaard, H.K. Jacobsen, Changes in hourly electricity consumption profiles and price elasticities in Denmark 2019–2022, *Energy* 308 (2024) 132798, <https://doi.org/10.1016/j.energy.2024.132798>, <https://www.sciencedirect.com/science/article/pii/S0360544224025726>.
- [7] M. Thoma, Electrical energy usage over the business cycle, *Energy Econ.* 26 (2004) 463–485, <https://doi.org/10.1016/j.eneco.2004.04.006>, <https://www.sciencedirect.com/science/article/pii/S0140988304000167>.
- [8] X. Wang, C. Si, J. Gu, G. Liu, W. Liu, J. Qiu, J. Zhao, Electricity-consumption data reveals the economic impact and industry recovery during the pandemic, *Sci. Rep.* 11 (2021) 19960, <https://doi.org/10.1038/s41598-021-98259-3>, <https://www.nature.com/articles/s41598-021-98259-3>.
- [9] M. Krarti, M. Aldubyan, Review analysis of COVID-19 impact on electricity demand for residential buildings, *Renew. Sustain. Energy Rev.* 143 (2021) 110888, <https://doi.org/10.1016/j.rser.2021.110888>, <https://www.sciencedirect.com/science/article/pii/S1364032121001829>.
- [10] V. van Zoest, K. Lindberg, F. El Gohary, C. Bartusch, Evaluating the effects of the COVID-19 pandemic on electricity consumption patterns in the residential, public, commercial and industrial sectors in Sweden, *Energy AI* 14 (2023) 100298, <https://doi.org/10.1016/j.egyai.2023.100298>, <https://www.sciencedirect.com/science/article/pii/S2666546823000708>.
- [11] S. Zhang, Q. Guo, R. Smyth, Y. Yao, Extreme temperatures and residential electricity consumption: evidence from Chinese households, *Energy Econ.* 107 (2022) 105890, <https://doi.org/10.1016/j.eneco.2022.105890>, <https://www.sciencedirect.com/science/article/pii/S014098832200072X>.
- [12] A. Gerossier, T. Barbier, R. Girard, A novel method for decomposing electricity feeder load into elementary profiles from customer information, *Appl. Energy* 203 (2017) 752–760, <https://doi.org/10.1016/j.apenergy.2017.06.096>, <https://www.sciencedirect.com/science/article/pii/S0306261917308486>.
- [13] J. Kraft, A. Kraft, On the relationship between energy and GNP, *J. Energy Dev.* 3 (1978) 401–403, <https://www.jstor.org/stable/24806805>.
- [14] F.M. Andersen, H.V. Larsen, R.B. Gaardestrup, Long term forecasting of hourly electricity consumption in local areas in Denmark, *Appl. Energy* 110 (2013) 147–162, <https://doi.org/10.1016/j.apenergy.2013.04.046>, <https://www.sciencedirect.com/science/article/pii/S0306261913003413>.
- [15] F.M. Andersen, H.V. Larsen, T.K. Boomsma, Long-term forecasting of hourly electricity load: identification of consumption profiles and segmentation of customers, *Energy Convers. Manag.* 68 (2013) 244–252, <https://doi.org/10.1016/j.enconman.2013.01.018>, <https://www.sciencedirect.com/science/article/pii/S0196890413000381>.
- [16] A. Diaz-Iglesias, X. Belauzarán, A.M. Florez-Tapia, Short-Term Power Demand Forecasting for Diverse Consumer Types to Enhance Grid Planning and Synchronisation, 2025. [arXiv:2506.04294 \[cs\]](https://arxiv.org/abs/2506.04294).
- [17] G. Pagliarini, C. Bonfiglioli, P. Vocale, Outdoor temperature sensitivity of electricity consumption for space heating and cooling: an application to the city of Milan, north of Italy, *Energy Build.* 204 (2019) 109512, <https://doi.org/10.1016/j.enbuild.2019.109512>, <https://www.sciencedirect.com/science/article/pii/S0378778819317360>.
- [18] F.M. Andersen, H.V. Larsen, N. Juul, R.B. Gaardestrup, Differentiated long term projections of the hourly electricity consumption in local areas. The case of Denmark West, *Appl. Energy* 135 (2014) 523–538, <https://doi.org/10.1016/j.apenergy.2014.08.075>, <https://www.sciencedirect.com/science/article/pii/S0306261914008861>.
- [19] D.D. Lee, H.S. Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* 401 (1999) 788–791, <https://doi.org/10.1038/44565>, <https://www.nature.com/articles/44565>, number: 6755.
- [20] D.D. Lee, H.S. Seung, Algorithms for non-negative matrix factorization, in: *Proceedings of the 13th International Conference on Neural Information Processing Systems, NIPS'00*, MIT Press, Cambridge, MA, USA, 2000, pp. 535–541.
- [21] J.V. Paatero, P.D. Lund, A model for generating household electricity load profiles, *Int. J. Energy Res.* 30 (2006) 273–290, <https://doi.org/10.1002/er.1136>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/er.1136>, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/er.1136>.
- [22] T. Räsänen, D. Voukantsis, H. Niska, K. Karatzas, M. Kolehmainen, Data-based method for creating electricity use load profiles using large amount of customer-specific hourly measured electricity use data, *Appl. Energy* 87 (2010) 3538–3545, <https://doi.org/10.1016/j.apenergy.2010.05.015>, <https://www.sciencedirect.com/science/article/pii/S0306261910001832>.
- [23] T. Teeraratkul, D. o'Neill, S. Lall, Shape-Based approach to household electric load curve clustering and prediction, *IEEE Trans. Smart Grid* 9 (2018) 5196–5206, <https://doi.org/10.1109/TSG.2017.2683461>, <https://ieeexplore.ieee.org/document/7880560>.
- [24] F.M. Andersen, P.A. Gunkel, H.K. Jacobsen, L. Kitzing, Residential electricity consumption and household characteristics: an econometric analysis of danish smart-meter data, *Energy Econ.* 100 (2021) 105341, <https://doi.org/10.1016/j.eneco.2021.105341>, <https://www.sciencedirect.com/science/article/pii/S0140988321002474>.
- [25] V. Michalakopoulos, E. Sarmas, I. Papias, P. Skaloumpakas, V. Marinakis, H. Doukas, A machine learning-based framework for clustering residential electricity load profiles to enhance demand response programs, *Appl. Energy* 361 (2024) 122943, <https://doi.org/10.1016/j.apenergy.2024.122943>, <https://www.sciencedirect.com/science/article/pii/S030626192400326X>.
- [26] A.S. Jin, K.T. Sanders, Characterizing residential sector load curves from smart meter datasets, *Appl. Energy* 366 (2024) 123316, <https://doi.org/10.1016/j.apenergy.2024.123316>, <https://www.sciencedirect.com/science/article/pii/S0306261924006998>.
- [27] J.A. Jardini, C.M.V. Tahan, M.R. Gouvea, S.U. Ahn, F.M. Figueiredo, Daily load profiles for residential, commercial and industrial low voltage consumers, *IEEE Trans. Power Deliv.* 15 (2000) 375–380, <https://doi.org/10.1109/61.847276>, <https://ieeexplore.ieee.org/document/847276>.
- [28] K. Bellinguer, R. Girard, A. Bocquet, A. Chevalier, ELMAS: a one-year dataset of hourly electrical load profiles from 424 French industrial and tertiary sectors, *Sci. Data* 10 (2023) 686, <https://doi.org/10.1038/s41597-023-02542-z>, <https://www.nature.com/articles/s41597-023-02542-z>.
- [29] C. Behm, L. Nolting, A. Praktikio, How to model european electricity load profiles using artificial neural networks, *Appl. Energy* 277 (2020) 115564, <https://doi.org/10.1016/j.apenergy.2020.115564>, <https://www.sciencedirect.com/science/article/pii/S030626192031076X>.
- [30] G.W. Hart, Nonintrusive appliance load monitoring, *Proc. IEEE* 80 (1992) 1870–1891, <https://doi.org/10.1109/5.192069>.
- [31] P.A. Schirmer, I. Mporas, Non-Intrusive load monitoring: a review, *IEEE Trans. Smart Grid* 14 (2023) 769–784, <https://doi.org/10.1109/TSG.2022.3189598>, <https://ieeexplore.ieee.org/abstract/document/9820770>.
- [32] Y. Zhu, S. Lu, Load profile disaggregation by blind source separation: a wavelets-assisted independent component analysis approach, in: *2014 IEEE PES General Meeting Textbook Conference & Exposition, 2014*, pp. 1–5, <https://doi.org/10.1109/PESGM.2014.6938947>, <https://ieeexplore.ieee.org/document/6938947>.
- [33] S. Semwal, D. Joshi, R.S. Prasad, D. Raveendhra, The practicability of ICA in home appliances load profile separation using current signature: a preliminary study, in: *2013 International Conference on Power, Energy and Control (ICPEC), 2013*, pp. 756–759, <https://doi.org/10.1109/ICPEC.2013.6527756>, <https://ieeexplore.ieee.org/document/6527756>.

- [34] G.M. Herath, M.H. Liyanage, C.J. Angammana, Assessment of Non-Intrusive load monitoring as a blind source separation problem, in: 2023 IEEE PES GTD International Conference and Exposition (GTD), 2023, pp. 156–160, <https://doi.org/10.1109/GTD49768.2023.00056>, <https://ieeexplore.ieee.org/document/10261743>.
- [35] M. Figueiredo, B. Ribeiro, A. de Almeida, Electrical signal source separation via nonnegative tensor factorization using on site measurements in a smart home, *IEEE Trans. Instrum. Meas.* 63 (2014) 364–373, <https://doi.org/10.1109/TIM.2013.2278596>, <https://ieeexplore.ieee.org/document/6588559>.
- [36] M. Matsumoto, Y. Fujimoto, Y. Hayashi, Energy disaggregation based on Semi-Binary NMF, in: P. Perner (Ed.), *Machine Learning and Data Mining in Pattern Recognition*, Springer International Publishing, Cham, 2016, pp. 401–414, https://doi.org/10.1007/978-3-319-41920-6_31.
- [37] A. Rahimpour, H. Qi, D. Fugate, T. Kuruganti, Non-Intrusive energy disaggregation using Non-Negative matrix factorization with sum-to-k constraint, *IEEE Trans. Power Syst.* 32 (2017) 4430–4441, <https://doi.org/10.1109/TPWRS.2017.2660246>, <https://ieeexplore.ieee.org/document/7835299>.
- [38] A. Miyasawa, Y. Fujimoto, Y. Hayashi, Energy disaggregation based on smart metering data via semi-binary nonnegative matrix factorization, *Energy Build.* 183 (2019) 547–558, <https://doi.org/10.1016/j.enbuild.2018.10.030>, <https://www.sciencedirect.com/science/article/pii/S0378778818302226>.
- [39] A.K. Zarabie, S. Das, An L0-Norm Constrained Non-Negative Matrix Factorization Algorithm for the Simultaneous Disaggregation of Fixed and Shiftable Loads, 2019, arXiv:1908.00142 [cs, eess].
- [40] S. Singh, A. Majumdar, Analysis Co-Sparse coding for energy disaggregation, *IEEE Trans. Smart Grid* 10 (2019) 462–470, <https://doi.org/10.1109/TSG.2017.2743763>, <https://ieeexplore.ieee.org/document/8016423>.
- [41] D. Piga, A. Cominola, M. Giuliani, A. Castelletti, A.E. Rizzoli, Sparse optimization for automated energy end use disaggregation, *IEEE Trans. Control Syst. Technol.* 24 (2016) 1044–1051, <https://doi.org/10.1109/TCST.2015.2476777>, <https://ieeexplore.ieee.org/document/7274721>.
- [42] J. Kolter, S. Batra, A. Ng, Energy disaggregation via discriminative sparse coding, in: *Proceedings of the 24th International Conference on Neural Information Processing Systems - Volume 1*, vol. 23, Curran Associates, Inc., 2010, <https://doi.org/10.5555/2997189.2997318>, https://papers.nips.cc/paper_files/paper/2010/hash/7810ccd41bf26faaa2c4e1f20db70a71-Abstract.html.
- [43] M. Wytock, J. Kolter, Contextually supervised source separation with application to energy disaggregation, *Proc. AAAI Conf. Artif. Intell.* 28 (2014), <https://doi.org/10.1609/aaai.v28i1.8769>, <https://ojs.aaai.org/index.php/AAAI/article/view/8769>, number: 1.
- [44] C. Dampyrou, A. Goichon, M. Ghienne, V. Tschannen, S. Schaack, Unsupervised separation of the thermosensitive contribution in the power consumption at a country scale, *Appl. Energy* 363 (2024) 123097, <https://doi.org/10.1016/j.apenergy.2024.123097>, <https://www.sciencedirect.com/science/article/pii/S030626192400480X>.
- [45] N. Gillis, *Nonnegative Matrix Factorization*, Data Science, Society for Industrial and Applied Mathematics, 2020, <https://doi.org/10.1137/1.9781611976410>, <https://epubs.siam.org/doi/book/10.1137/1.9781611976410>.
- [46] Y.-T. Guo, Q.-Q. Li, C.-S. Liang, The rise of nonnegative matrix factorization: algorithms and applications, *Inf. Syst.* 123 (2024) 102379, <https://doi.org/10.1016/j.is.2024.102379>, <https://www.sciencedirect.com/science/article/pii/S0306437924000371>.
- [47] Y. Liu, R. Jin, L. Yang, Semi-supervised multi-label learning by constrained nonnegative matrix factorization, in: *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1, AAAI'06*, AAAI Press, Boston, Massachusetts, 2006, pp. 421–426.
- [48] H. Liu, Z. Wu, X. Li, D. Cai, T.S. Huang, Constrained nonnegative matrix factorization for image representation, *IEEE Trans. Pattern Anal. Mach. Intell.* 34 (2012) 1299–1311, <https://doi.org/10.1109/TPAMI.2011.217>, <https://ieeexplore.ieee.org/document/6072214>.
- [49] V.P. Pauca, J. Piper, R.J. Plemmons, Nonnegative matrix factorization for spectral data analysis, *Linear Algebra Appl.* 416 (2006) 29–47, <https://doi.org/10.1016/j.laa.2005.06.025>, <https://www.sciencedirect.com/science/article/pii/S002437950500340X>.
- [50] S. Jia, Y. Qian, Constrained nonnegative matrix factorization for hyperspectral unmixing, *IEEE Trans. Geosci. Remote Sens.* 47 (2009) 161–173, <https://doi.org/10.1109/TGRS.2008.2002882>, <https://ieeexplore.ieee.org/document/4694061/>.
- [51] X. Liu, W. Xia, B. Wang, L. Zhang, An approach based on constrained nonnegative matrix factorization to unmix hyperspectral data, *IEEE Trans. Geosci. Remote Sens.* 49 (2011) 757–772, <https://doi.org/10.1109/TGRS.2010.2068053>, <https://ieeexplore.ieee.org/document/5585746/>.
- [52] Z. Chen, A. Cichocki, *Nonnegative Matrix Factorization with Temporal Smoothness and/or Spatial Decorrelation Constraints*, 2004.
- [53] M. Hosseinzadeh Aghdam, A novel constrained non-negative matrix factorization method based on users and items pairwise relationship for recommender systems, *Expert Syst. Appl.* 195 (2022) 116593, <https://doi.org/10.1016/j.eswa.2022.116593>, <https://www.sciencedirect.com/science/article/pii/S0957417422000896>.
- [54] C. Peng, K.-C. Wong, A. Rockwood, X. Zhang, J. Jiang, D. Keyes, Multiplicative algorithms for constrained non-negative matrix factorization, in: *2012 IEEE 12th International Conference on Data Mining, IEEE, Brussels, Belgium, 2012*, pp. 1068–1073, <https://doi.org/10.1109/ICDM.2012.106>, <http://ieeexplore.ieee.org/document/6413807/>.
- [55] Terna, Total load, 2024. <https://dati.terna.it/en/load#total-load>.
- [56] Terna, Consumi 20XX, Technical Report, Terna S.p.A., 2024. <https://www.terna.it/it/sistema-elettrico/statistiche/publicazioni-statistiche>.
- [57] Terna, Enel Household Consumption Index, 2024.
- [58] Terna, IMCEI industrial consumption, 2024. <https://dati.terna.it/en/load/imcei>.
- [59] Terna, IMSER services consumption, 2024. <https://dati.terna.it/en/load/imser>.
- [60] J.O. Ramsay, B.W. Silverman, *Functional Data Analysis*, Springer Series in Statistics, Springer, New York, NY, 2005, <https://doi.org/10.1007/b98888>, <http://link.springer.com/10.1007/b98888>.
- [61] R.J. Hyndman, G. Athanasopoulos (3rded), *Forecasting: Principles and Practice*, OTexts, Melbourne, Australia, 2021, <https://otexts.com/fpp3/>.