



Universidad de Oviedo
Universidá d'Uviéu
University of Oviedo

Departamento de Ingeniería Eléctrica, Electrónica,
de Computadores y Sistemas

PHD THESIS

**Energy disaggregation techniques for visualization and
improvement of energy efficiency in processes and buildings**

PROGRAMA DE DOCTORADO EN ENERGÍA Y CONTROL DE PROCESOS.

Autor: Diego García Pérez
Directores: Ignacio Díaz Blanco
Manuel Domínguez González

MAYO 2021



Universidad de Oviedo
Universidá d'Uviéu
University of Oviedo

Departamento de Ingeniería Eléctrica, Electrónica,
de Computadores y Sistemas

PHD THESIS

**Energy disaggregation techniques for visualization and
improvement of energy efficiency in processes and buildings**

TESIS DOCTORAL

**Técnicas de desagregación de energía para la visualización y
mejora de la eficiencia energética en procesos y edificios**

Memoria presentada para la obtención del grado de Doctor
por la Universidad de Oviedo, en conformidad con los requisitos
del PROGRAMA DE DOCTORADO EN ENERGÍA Y CONTROL DE PROCESOS.

Autor: Diego García Pérez
Directores: Ignacio Díaz Blanco
Manuel Domínguez González

GIJÓN, MAYO 2021

RESUMEN

A lrededor del 45% de la energía consumida en los países desarrollados se destina a hogares, servicios públicos y actividades comerciales. Por esta razón, la mejora de la eficiencia energética en estas instalaciones ha cobrado mucha importancia en las últimas décadas. Este reciente interés en la mejora de la eficiencia energética ha coincidido en el tiempo con el auge de los modelos basados en datos, tales como el *aprendizaje automático* o la *visualización de datos*, lo que ha provocado que muchas investigaciones en eficiencia energética integren soluciones basadas en datos, con el fin último de incrementar la conciencia energética del usuario.

Un ejemplo de estos enfoques basados en datos es el análisis *Non-Intrusive Load Monitoring* (NILM), capaz de proporcionar a los usuarios una representación basada en partes de la demanda energética, estimando cada consumo individual utilizando únicamente el consumo global de la instalación como entrada.

A pesar de la gran cantidad de soluciones propuestas para el problema NILM a lo largo de las últimas décadas, solo unas pocas han abordado el problema NILM en grandes edificios no residenciales. Además, los modelos NILM dentro de la literatura suelen ser valorados por la precisión de la descomposición obtenida, en vez de por la calidad de la información recibida por los usuarios y la mejora de su conciencia energética. Teniendo esto en cuenta, la presente tesis tiene dos objetivos principales: 1) explorar el problema NILM en grandes edificios no residenciales y cómo adaptar los algoritmos NILM existentes a este escenario y 2) mejorar la información proporcionada a los usuarios una vez obtenida la desagregación NILM, integrándose en visualizaciones interactivas para formar modelos de analítica visual.

Los modelos de analítica visual explotan las sinergias entre el aprendizaje automático, la visualización de datos y los mecanismos de interacción para incorporar al usuario y su experiencia previa a un bucle infinito de análisis, en el que el conocimiento sobre el problema bajo estudio crece y se asienta de forma significativa. Este paradigma es adaptado al problema NILM, definiendo nuevas vías de interacción con los algoritmos de desagregación de energía y con la instalación, por medio de las cuales el usuario puede configurar el análisis de acuerdo con su conocimiento previo y modificar su comportamiento con la instalación para mejorar la eficiencia energética de la misma.

Nuestro enfoque de analítica visual en NILM es abordado mediante la integración de una descomposición no supervisada resultante de aplicar el algoritmo *Non-negative Matrix Factorization* (NMF) en una estructura tipo *cubo de datos*, que nos permite no solo agilizar las interacciones de filtrado y selección, sino que también nos proporciona valores de salida apropiados para ser tratados con técnicas de visualización de datos en una y dos dimensiones. Se utilizan datos reales procedentes del Hospital de León para demostrar que las visualizaciones de la descomposición NMF y los mecanismos de interacción fluida soportados por el cubo de datos permiten al usuario establecer visualmente correlaciones entre los patrones latentes obtenidos de NMF y subsistemas críticos en grandes edificios, como los sistemas de confort térmico.

En busca de una desagregación más detallada, se exploran modelos NILM novedosos basados en *redes neuronales profundas* (DNN, según su denominación en inglés: deep neural networks), los cuales se adaptan a nuestro paradigma de análisis visual, así como a la naturaleza de los grandes edificios. Para ello, se propone un nuevo modelo de *denoising Auto-Encoder* totalmente convolucional (FCN-dAE) apto para NILM en grandes edificios, que supera a los enfoques de NILM más avanzados en los datos del hospital, y muestra una mejor eficiencia computacional que los modelos anteriores en términos de número de pesos y tiempo de inferencia.

Para cerrar el bucle de análisis entre las técnicas NILM basadas en DNN y el usuario, y siguiendo nuestra formulación de analítica visual, se introducen en nuestro enfoque FCN-dAE vías de interacción fluidas basadas en *mecanismos de condicionamiento*. En concreto, nuestro FCN-dAE se modula mediante la técnica de condicionamiento *Feature-wise Linear Modulation*, la cual permite añadir una entrada auxiliar al modelo, mediante la cual el usuario puede modular el comportamiento de la red neuronal según sus intenciones. Esto conduce a transiciones fluidas y continuas entre los consumos individuales estimados, que no sólo hacen que el modelo sea adecuado para los enfoques de analítica visual, sino que también pueden ayudar a comprender los mecanismos internos de las redes neuronales profundas utilizadas.

ABSTRACT

About 45% of the energy demand in developed countries is consumed in households and in public and commercial services. For this reason, the improvement of electrical energy efficiency in facilities has attracted much attention in recent decades. This recent pursuit of energy efficiency has coincided in time with the rise of data-driven models, such as *machine learning* or *data visualization* techniques, which has led many researches to incorporate data-driven solutions into energy efficiency problems, with the ultimate goal of increasing the energy awareness of the user.

An example of these data-driven approaches is the *Non-Intrusive Load Monitoring* (NILM) approach, which provides users with a parts-based representation of the energy demand by estimating each individual consumption, using only the energy demand measurements of the entire facility as input.

Despite the huge amount of solutions proposed to the NILM problem over the last decades, only a few have addressed the NILM problem in large non-residential buildings. Furthermore, NILM models in the literature are ranked by the accuracy of the resulting energy disaggregation, rather than by the quality of the information received by users and the improvement of their energy awareness. Bearing this in mind, the present PhD thesis has two main goals: 1) exploring the NILM problem in large non-residential buildings and how to adapt existing NILM algorithms to this scenario; and 2) improving the user feedback obtained from NILM algorithms by integrating the resulting disaggregation into interactive visualizations, following the *visual analytics* paradigm.

Visual analytics approaches exploit the synergies between machine learning, data visualization and interaction mechanisms to incorporate the user and its prior knowledge into an infinite loop of analysis, in which his understanding about the problem under study is significantly boosted. This paradigm is adapted to the NILM problem, defining new interaction paths with the energy disaggregation algorithms and with the facility,

whereby the user can tweak the analysis according to his prior knowledge and modify his behavior with the facility in order to improve the overall energy efficiency.

In a first approach, the problem of visual analysis in NILM is addressed by arranging the unsupervised decomposition obtained from the *Non-negative Matrix Factorization* (NMF) algorithm in a powerful *OLAP data cube* structure, which not only speeds up filtering and selection interactions, but also provides output values suitable for 1D and 2D insightful data visualization techniques. Real data from the Hospital of León (Spain) are used to show that the visualizations of the NMF decomposition and the fluid interaction pathways supported by the data cube allow the user to visually establish correlations between the obtained latent patterns and important subsystems in large buildings, such as thermal comfort systems.

Seeking a more detailed disaggregation, novel *deep neural networks* (DNN) models are explored and adapted to our visual analytics paradigm and the nature of large buildings. A novel *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE) model is proposed for NILM in large buildings, which outperforms the state-of-art NILM approaches on data from the hospital, and shows better computational efficiency than the previous models in terms of number trainable weights and inference time.

In order to close the loop of analysis between DNN-based NILM techniques and the user, following our visual analytics formulation, fluid interaction pathways based on *conditioning mechanisms* are introduced into our FCN-dAE approach. More specifically, our FCN-dAE is conditioned by the general-purpose *Feature-wise Linear Modulation* conditioning mechanism to add an auxiliary input, by which the user can modulate the behavior of the model according to his intentions. This leads to fluid and continuous transitions in the estimated individual consumptions that not only make the model suitable for visual analytics approaches, but may also help to understand the inner mechanisms of the DNN-based approaches.

CONTENTS

	Page
List of Tables	ix
List of Figures	xi
1 Introduction	1
1.1 Introduction	1
1.2 Problem Formulation	5
1.3 Objectives	12
1.4 Structure of the document	13
2 NILM framework and disaggregation techniques	15
2.1 NILM framework	15
2.1.1 Data acquisition	16
2.1.2 Data preparation and feature extraction	18
2.1.3 Load inference and learning	22
2.1.4 Disaggregation and deployment	24
2.2 Supervised disaggregation algorithms	25
2.2.1 Optimization methods	25
2.2.2 Pattern recognition methods	26
2.2.3 Deep learning methods	28
2.3 Unsupervised disaggregation algorithms	29
2.3.1 Unsupervised HMM-based models	30
2.3.2 Source Separation methods	30
2.3.3 Unsupervised graph-based signal processing	31
2.4 Evaluation of NILM methods	31
2.5 Challenges and open issues	36

3	Visual analytics for NILM in large buildings	39
3.1	Differences between residential and non-residential facilities	39
3.1.1	Datasets	40
3.1.2	Nature of the consumptions in residential and non-residential facilities	44
3.1.3	Visual analytics as a first approach to NILM in large building . . .	51
3.2	Blind Source Separation approaches for NILM in large buildings	53
3.2.1	Non-negative Matrix Factorization formulation	54
3.2.2	Complexity of NMF	55
3.2.3	Objective function and SVD-based initialization	56
3.2.4	Sparseness of NMF decomposition	57
3.2.5	Adaptation of NMF to time series	59
3.3	Visual techniques in large buildings NILM	60
3.3.1	Principles of visual coding	61
3.3.2	Principles of interaction	65
3.3.3	The data cube paradigm	67
3.3.4	Data visualization techniques	73
3.4	Experiments	76
3.4.1	Experimental set-up	77
3.4.2	Analysis and interpretation of NMF components	80
3.4.3	Evaluation of the visual analytics application through use cases .	83
3.5	Conclusions	91
4	NILM approaches for large buildings based on Deep Learning techniques	95
4.1	Related work	97
4.2	Denoising Auto-Encoder models as NILM systems in large buildings . . .	103
4.2.1	Vanilla denoising Auto-Encoders (dAE)	103
4.2.2	Disaggregation process and the input window size selection	107
4.2.3	Fully-convolutional denoising auto-encoder (FCN-dAE)	109
4.3	Integration of FCN-dAE into visual analytics	112
4.3.1	Conditioning in deep neural networks models	113
4.3.2	Multi-task FCN-dAE	117
4.4	Results	119
4.4.1	Experimental set-up	119

4.4.2	Input sequence length effects	121
4.4.3	Performance of FCN-dAE in residential facilities	127
4.4.4	Computational efficiency of FCN-dAE	128
4.4.5	Use cases in large buildings	130
4.4.6	Performance of multi-task FCN-dAE	138
4.4.7	Interpretation of conditioning mechanisms	141
4.4.8	Conclusions	143
5	Conclusions and future work	145
5.1	Discussion and final conclusions	145
5.2	Thesis contributions	148
5.3	Future work	149
6	Conclusiones y trabajo futuro	151
6.1	Discusión y conclusiones finales	151
6.2	Contribuciones de la tesis	154
6.3	Trabajo futuro	156
	Agradecimientos	159
	Acknowledgments	161
A	Publications	163
A.1	Principal publications	164
A.2	Related publications	167
	Bibliography	169

LIST OF TABLES

TABLE	Page
2.1 Common binary classification metrics for NILM.	33
2.2 Common energy estimations metrics for NILM. Note that the first rows of the table indicate the notation used in the metrics.	35
3.1 General features of UK-DALE dataset extracted from [1].	41
3.2 Meters used in the hospital’s submetering network.	42
3.3 General details of the CGBT2.Red-Grupo subset from the <i>Hospital de León dataset</i>	44
3.4 Variables used as <i>context factors</i> in visual analytics approach. The outcome of the NMF decomposition will be added to these variables before creating the data cube.	78
3.5 Pearson coefficients computed from the raw samples of the initial table, and between the 3-th, 8-th components, the individual consumptions of the cooling machines P1 and T2, and the ambient temperature conditioned to low (L) [100-200], medium (M) [100-200], high (H) [100-340] values of the component, as well as the whole-range correlation (WR) [0-350].	89
3.6 Group-based Pearson coefficients computed from the data cube output, and between the 3-th, 8-th components, the individual consumptions of the cooling machines P1 and T2, and the ambient temperature conditioned to low (L) [100-200], medium (M) [100-200], high (H) [100-340] values of the component, as well as the whole-range correlation (WR) [0-350].	89
4.1 Disaggregation performance of NILM models with different input sequence lengths in all individual nodes from Hospital. In the first column a FHMM approach is included in order to compare the DNN-based results with a HMM-based approach. The NEP, RMSE and MAE metrics are used as evaluation metrics.	122

4.2 Disaggregation performance of models with different input sequence lengths in all individual nodes from Hospital. In the first column a FHMM approach is included in order to compare the DNN-based results with a HMM-based approach. The NEP, RMSE and MAE metrics are used as evaluation metrics. 127

4.3 Comparison of disaggregation performance between sequence-to-sequence models and *multi-task FCN-dAE*. 139

LIST OF FIGURES

FIGURE	Page
1.1 Final energy consumption by sector in million tones of oil since 2015 to 2018. Source [2].	5
1.2 Types of appliances in residential facilities.	7
1.3 Two-way analysis NILM-User.	9
1.4 Concept model of getting actionable feedback from a NILM system. Inspired by [3]	10
2.1 CRISP-DM steps defined in the original article [4].	16
2.2 Diagram of NILM framework.	17
2.3 Event-based energy disaggregation process. After the event detection, each event $e^{(i)}$ is identified by means of features computed on the <i>steady state</i> or <i>transient</i> sequences.	20
2.4 DNN-based energy disaggregation process. The main consumption of the facility is split in a set of overlapped sequences, using a windowing operation. The resulting windows are processed by a trained DNN model to produce a set of overlapped windows of the estimated individual consumption. The final individual consumption waveform is computed by aggregating out windows of DNN model by means of an aggregation function, such the mean.	23
2.5 Taxonomy of load disaggregation methods.	24
2.6 Conceptual diagram of pattern recognition NILM models.	27
3.1 Submetering architecture of the Hospital de León dataset. The main circuits are represented as a vertical line, and the important connections are depicted as a point. All the measured connections are highlighted and identified by its label and the type of meter installed.	43
3.2 Average energy demand of the node CGBT-2.Red-Grupo grouped by different time scales (months, days of week and hours).	45

3.3	Average energy demand of House 1 from UK-DALE repository grouped by different time scales (months, days of week and hours).	46
3.4	Daily profile of all the houses from UK-DALE repository.	46
3.5	Histograms of main consumption of CGBT-2.Red-Grupo and House 1 from UK-DALE.	47
3.6	Intuition about the individual consumptions of the node CGBT-2.Red-Grupo. a) Energy demand of the individual nodes and b) histograms of the individual nodes.	49
3.7	Intuition about the individual consumptions of the House 1 from UK-DALE repository. a) Energy demand of the individual appliances and b) histograms of the appliance-specific consumptions. Note that these histograms are in logarithmic scale and most of the recorded samples are zero.	50
3.8	Diagram of our visual analytics perspective of NILM particularized for large buildings.	52
3.9	NMF decomposition of faces from [5]. Matrix on the left represents the matrix of bases \mathbf{W} and the matrix on the right represents the matrix of activations \mathbf{H}	56
3.10	Intuition about the coordinate gradient descent of the NMF decomposition. Possible NMF solutions rely on the feasible positive regions.	58
3.11	Example of <i>small multiples</i> from (https://pewrsr.ch/3o4Suml).	62
3.12	Data type taxonomy extracted from [6].	63
3.13	Combination of marks and channels.	63
3.14	Taxonomy of channels effectiveness. Inspired by the discussion about the use of channels in [6].	64
3.15	Sense making process extracted from [7].	66
3.16	Intuition of data cube formation and terminology.	69
3.17	Examples of data cube operations. In the first row the <i>projection operation</i> is shown by firstly projecting the initial cube to a 2D cube and then to 1D cube. In the second row a combination of <i>selection</i> , <i>projection</i> and <i>aggregation</i> operations is shown.	71
3.18	Interactive barchart technique for representing 1D projections of the data cube.	75
3.19	Heatmaps which represent the 2D projections of the data cube. a) represents the average of the total power demand of the hospital (PotenciaPF) grouped by hour and day of week. b) represents the average of the total power demand grouped by days of year arranged in a calendar distribution.	76

3.20	Snapshot of sparklines of NMF basis consumptions.	77
3.21	General view of the suggested VA application.	80
3.22	Dialog window which enables the user to interactively choose the parameters of the NMF decomposition.	80
3.23	Reconstruction error according to the number of components.	81
3.24	Some components from the trained NMF model with $L = 20$. The first three factors are more relevant both in daily profile and throughout the year than the last factors.	82
3.25	Filtering the largest values of the first three components by means of drag gestures in their associated barcharts. a) Weekday patterns are spotted in both the calendar and weekly heatmaps after filtering by the largest values of the first component and selecting count as the aggregation function. In b), the same analysis as in a) shows that the highest values of the second component highlight the weekends and holidays. The calendar in c) shows that the largest values of the third component are related to a seasonal behavior.	85
3.26	Configuration of the interactive data cube whereby an insightful visual correlation was found in the third and eighth components. In a), correlations can be established between the third component, the ambient temperature and the consumptions of the cooling machines of type T. In the same way, b) shows how the low and medium values of the 8-th component can be strongly related to low temperatures and slightly related to the consumption of the P1 machine.	86
3.27	The exploration of the 18th component as an example of analysis of a non-periodic component. The user focus the analysis on one of the highlighted days, looking for evidences in the rest of the attributes that can explain the causation behind the component.	90
3.28	Searching for the causes of the anomalies. After selecting the 2nd of May, an anomaly in the starting of the cooling machines is discovered by means of filters in the individual consumptions of the cooling machines.	92
4.1	FHMM diagram extracted from [8].	99
4.2	Intuition of sequence-to-sequence DNN-based NILM models.	101
4.3	The studied dAE-based topologies. a) Vanilla dAE architecture where each cell symbolizes a layer explained by the number of units in the case of fully-connected layers; or the kernel size, stride and number of kernels in the case of convolutional layers. b) Proposed FCN-dAE architecture made entirely by 1D convolutional layers and 1D transposed convolutional layer.	106

4.4	Disaggregation process for DNN-based approaches.	108
4.5	Different estimations of fridge consumption produced by a dAE model trained for different input sequence lengths.	109
4.6	Graphical representation of the convolutional layers, receptive field and the composition of convolutional layers.	111
4.7	Graphical diagrams of the idea of conditioning DNN models	114
4.8	FiLM conditioning approach.	116
4.9	Multi-task FCN-dAE architecture.	118
4.10	Estimations for the node <i>Radiologia1</i> obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.	123
4.11	Estimations for the node <i>RehabilitacionB</i> obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.	124
4.12	Power spectral density of the residuals calculated from the disaggregations resulting from applying the dAE and FCN-dAE architectures, trained with input window sizes of 30, 120 and 1440 minutes, to the individual node <i>Radiologia1</i> . The highlighted harmonics correspond to the errors made by the models at daily (freq. = 0.041 hour^{-1}) and at 75 minutes (freq. = 0.8 hour^{-1}) periods.	125
4.13	Estimations for the node <i>RehabilitacionA</i> obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.	126
4.14	Estimations for the appliance <i>washing machine</i> (WM) obtained after applying all DNN-based architectures, trained with input window sizes of 30, 60, 90 and 120 minutes.	128
4.15	Estimations for the appliance <i>fridge</i> (FR) obtained after applying all DNN-based architectures, trained with input window sizes of 30, 60, 90 and 120 minutes.	129
4.16	Computational efficiency in terms of time of computation and number of weights per topology. The first plot depicts the total amount of trainable weights of all the suggested models. Second plot indicates the time in seconds required for each model to compute the estimation from one window of total consumption. Note that the axes in both graphs are logarithmic.	130

4.17	Example of gap filling when the failure only affects the individual consumption. This case is exemplified with the estimation of individual node <i>RehabilitationB</i> computed by dAE and FCN-dAE approaches with input sequence length of 1440 minutes.	132
4.18	Obtained RMSE evolution after applying FCN-dAE and dAE models to corrupted test main consumption segments with different gap lengths.	133
4.19	Example of gap filling when the failure also affects the main consumption of the facility. This case is exemplified with the estimation of individual node <i>RehabilitationB</i> computed by dAE and FCN-dAE approaches with input sequence length of 1440 minutes and for 9-hour gaps.	134
4.20	Explanatory diagram of the mechanism for obtaining the residuals from the real measured individual consumptions and the estimated individual consumptions obtained.	135
4.21	Sequence from node <i>RehabilitationB</i> with a novelty synthetically introduced (red color). In the second row, the residual $ p_m(t) - \hat{p}_m(t) $ between the estimated and measured individual consumption is shown. In this case the estimation $\hat{p}_m(t)$ is computed by means of a FCN-dAE model with an input sequence length of 1440 minutes.	136
4.22	Probably density of normal and abnormal residuals for all individual nodes computed by means of the estimations produced by the FCN-dAE with 1440 minutes input sequence length.	137
4.23	Precision-recall curves for residuals computed from all the proposed models ordered by window size. Grey lines indicate the precision-recall curve for a model with the specified F1-score.	138
4.24	Examples of individual nodes from the Hospital of León disaggregated by sequence-to-sequence DNN models and by suggested multi-task FCN-dAE. All models were trained with 1440-minutes input sequences.	140
4.25	Comparison of computational efficiency between suggested multi-task FCN-dAE and sequence-to-sequence DNN-based models.	141
4.26	Examples of transitions between two individual nodes collected at the output of our multi-task FCN-dAE when continuous trajectories between the one-hot encodings of the nodes are defined in the input S_D	142
4.27	UMAP projections of the latent space representation, after apply our multi FCN-dAE model to 300 sequences of main consumptions modulated by all the available individual nodes.	143

A.1 Article published in the *IEEE Transactions on Smart Grid* journal [9] with the results of chapter 4. 164

A.2 Article published in the *Energy and Buildings* journal [10] with the results of chapter 3. 165

A.3 Article published in the *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* [11]. It includes the idea of applying the NMF decomposition to energy-related data. 166

A.4 Article published in the *Oxford Bioinformatics* journal [12]. It is related to the idea of fluid reconfigurations of views according to the user hypotheses similar to the data cube interaction pathways studied in chapter 3. 167

A.5 Article published in the *Computers in Industry* journal [13]. A novel visual analytics solution boosted by a data cube structure similar to the presented in chapter 3 is suggested for the monitoring of a cold roll mill. 168

INTRODUCTION

In this first chapter, the motivation, the problem formulation and the main objectives established for this doctoral thesis will be presented. Finally, in the last part of the chapter, the structure of the rest of the manuscript will be detailed.

1.1 Introduction

Modern society is undoubtedly living in the era of digitalization, where each individual produces and consumes a large amount of data in the daily course of her or his activities. As an example, the simple fact of searching a restaurant involves a process where modern web search engines use previous information, such as users historical preferences, opinions or localizations, in order to suggest a personalized list of restaurants. The web searchers example is not an isolated case, but instead the use of algorithms that enhance the service quality by means of the data usage is widespread nowadays in many common areas, such as social networks, bank services, anomaly detection in industry, medical image analysis, etc., demonstrating that data have become a valuable resource. However, extracting useful knowledge from data is not straightforward, since the information gathered nowadays quite often satisfies the *big data* paradigm, which is defined by the rule of the 3 *v*'s [14]: the *volume* of the collected data has extraordinarily increased in the last decades due to the emergence of the *information and communications technology* (IT); the *velocity* with which the information is generated in modern applications is in

real-time or near real-time; and, last but not least, datasets are come from a wide variety of domains and sources such as image, time series, GPS signals, etc..

In this scenario, traditional analytics based on mathematical models or fixed rule-based computing algorithms cannot cope with modern data-driven problems, where the principles behind the data behavior are so complex that it is difficult to model them manually by these classic approaches. For this reason, computer science community have focused their efforts on developing more flexible techniques, which disentangle the data by learning patterns by themselves or by helping the user to intuitively extract hidden knowledge through efficient representations. The first category is called *machine learning* (ML) and it is defined as those algorithms that *learn* to solve a well-posed task from the experience gained in a *training* process [15], in which the algorithm is faced to examples of the task to be solved. On the other hand, the second category is based on *data visualization* (DV) techniques [6, 16] where it is the user who solves the task by analyzing visual and immersive representations of the data. DV-based techniques provide users with intuitive views that exploit the extraordinary cognitive abilities of human beings to extract patterns hidden on intricate datasets making the decision-making process easier.

Nowadays, ML models, especially *Deep Neural Networks* (DNN) considered within of the *Deep Learning* (DL) paradigm, have achieved extraordinary results in intricate tasks, such as image classification and segmentation [17, 18], natural language processing [19], finding new drug molecules [20], etc.. Although these models are often as accurate as human beings in solving complicated tasks (sometimes better), they are only suitable for well-defined problems, i.e., problems where the task to be solved, the metrics used for evaluating the performance and the training data can be perfectly identified. In contrast, when the problem is ill-posed because either the target problem is in a preliminary stage and there is no enough domain knowledge to set a specific learning task, or the context of the problem is continuously changing, human beings outperform ML algorithms, since we are capable of solving problems by means of an iterative and exploratory analysis. If the data are appropriately presented, humans may be able to integrate their previous experiences in similar problems into the analysis and apply their powerful cognitive system to find abstract connections behind the data. For this reason, *visual analytics* (VA) [21, 22] approaches which exploit the synergies between ML models, DV techniques and interaction mechanisms have attracted a lot of attention in the last decade. As a result of this growing interest, powerful visual tools where the user intuitively gains knowledge about the problem in an iterative human-machine process have emerged

[23, 24].

The aforementioned concepts can be transferred to the industrial sector, which may enable novel perspectives where the huge volume of information measured from processes could be exploited to reveal undiscovered behaviors of the systems, prevent catastrophic failures or improve their energy efficiency. The use of data in the industrial environment is previous to the appearance of the aforementioned data-based techniques, since many of its processes are regulated by feedback control or automation systems that, by definition, need detailed information about the condition of the system to be controlled. Despite the high availability of data in this sector, process condition variables measured from industrial systems or large buildings yield high-dimensional datasets which often suffer from well-known problems related to the data acquisition process, such as gaps or noise in the measurements.

In this unfriendly scenario, *finding valuable knowledge is not a trivial task and it often involves tackling the problem from different perspectives, splitting it into self-contained but interconnected subtasks*. Although these subproblems could be solved from the ML point of view, very often the domain knowledge of an expert may be crucial in the interpretation of the ML results and in correlating the subproblems between them. In this sense, *VA-based approaches can easily integrate the expert knowledge into the analysis*, enhancing the ML findings with visualizations that integrate context process variables and where, by means of the interaction mechanisms, an expert user can insightfully discover connections between the ML outcome and recognizable process' events. These correlations, which could not be obtained with ML approaches alone, nor with single visualizations, lead to new knowledge about the process and increase the confidence in the results.

A suitable industrial area of application for this approach is the energy efficiency field, where the consumption awareness of the user is crucial in order to reduce the waste of energy [25, 26]. A better understanding of how the energy is spent leads the user to improve his consumption profile, identifying and reducing inefficient behaviors. An intuitive way of explaining the energy consumption of a facility is to separately monitor the consumptions of all the plugged-in devices. For instance, in the case of an electrical household installation, decomposing the whole-house load into a sum of appliance-specific loads lets the user to identify how much energy has been spent in each device and when they have been turned on over the time. This *Load Monitoring*

approach (LM) is a parts-based representation that, if there is no energy generator among the connected devices, results in a positive sum of signals. This non-negative nature of the components suggests that they are the *parts of a whole* and the analysis of the efficiency can now be seen as getting insight by learning individually the resulting parts. Learning an object through its parts is an operation that has a strong connection with the human perception process as it is shown in [27, 28], and therefore it fits perfectly with the paradigm of visual analytics. The general hypothesis of this approach is that *the parts-based representation will result in an increase of the user's energy awareness and prompt actions for the improvement of the energy efficiency such as establishing better consumption schedules and detecting high consumptions or malfunctions in the individual consumptions.*

The easiest way to gather information from each plugged-in appliance is to install a meter per appliance, but this *Intrusive Load Monitoring* (ILM) approach would cause a great impact on the installation and in terms of set-up cost, making ILM methods undesirable for massive deployment. To solve this problem, Hart [29] proposed a cost-effective approach known as *Non-Intrusive Load Monitoring* (NILM), which aims to estimate individual consumptions, that are not directly measured, from the total consumption signal. This non-intrusive approach is able to disaggregate the energy consumption of an installation as the parts of a whole, without incorporating any element or additional cost to the installation. NILM techniques have attracted a lot of attention in the last decades, since they can help to reduce the energy waste in residential buildings and large commercial buildings, which are two of the highest energy demanding sectors in the European Union (see Figure 1.1), consuming approximately half of the total energy consumed in the member states.

This interest has led to an intense research in NILM, especially with the rise of the DL-based NILM techniques, and to the emergence of technology companies that have transferred the concept of NILM into a final product [30, 31]. In spite of all advances on NILM achieved during the past decades and the incipient deployment of NILM systems in real facilities, it is still an active research field with open questions, such as its scalability to large installations such as large public buildings, its integration into interactive applications, the automatic detection of the number of plugged-in devices or the transferability of NILM models to contexts other than those used for training (i.e., in facilities of the same nature but in other continent).

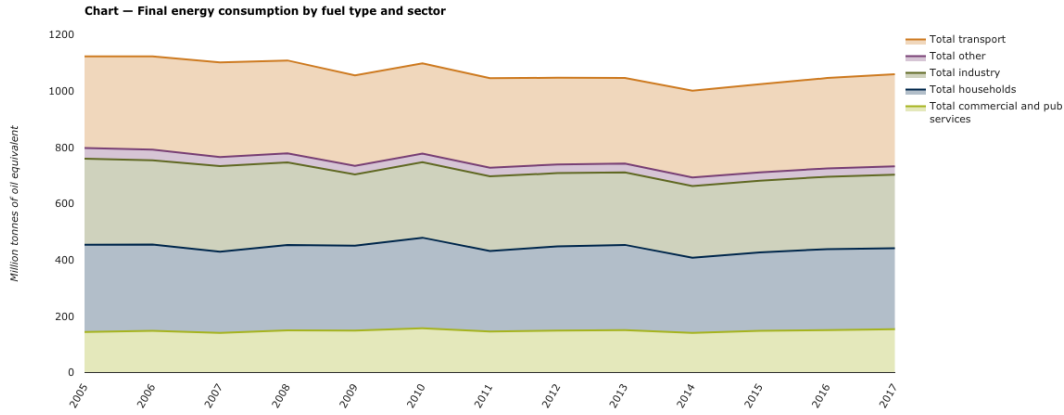


Figure 1.1: Final energy consumption by sector in million tonnes of oil since 2015 to 2018. Source [2].

Taking all these concepts into consideration, *the object of this doctoral thesis is to explore novel approaches of NILM that integrate elements of data visualization, machine learning algorithms and mechanisms of interaction into an intuitive presentation of the results with the aim of increasing the energy awareness. Moreover, the particularities of applying novel NILM models to large systems, such as public buildings, shall be studied in detail. The synergies between the application of novel NILM techniques to large systems and the powerful interpretability and usability provided by data visualization and interaction paves the way to unexplored approaches that can significantly enhance the decision-making process in the maintenance of a large installations.*

1.2 Problem Formulation

The idea of NILM was introduced in the 90s by Hart in [29], who defined for the first time a method capable of identifying the appliance loads from the total load measurements. Hence, each appliance-specific load is obtained from only one signal taken from the power supply connection without placing appliance-specific meters, as in previous ILM approaches. The problem is formulated as the decomposition of the whole energy load of the facility under study into a positive sum of individual loads. Thus, the whole-installation measurement, which can be any electrical variable such as active, reactive power or currents, is defined as $P(t)$ and it can be mathematically formulated as follows:

$$P(t) = p_1(t) + p_2(t) + \dots + p_m(t) \quad (1.1)$$

where the total electric measurement at time t is the sum of the n appliance-specific signals $p_i(t)$ corresponding to the downstream appliances. Note that if the network lacks generative systems that deliver power to the grid, the components p_i are always positive, resulting in a *parts-of-a-whole representation*. The nature of each appliance defines its specific load p_i and the whole-installation load $P(t)$ depends on the number of appliances which are switched-on, at time t . In order to model the number of switched-on devices, Hart defines a *switch process* vector $\mathbf{a}(t) \in \mathbb{R}^m$ which contains the state of the appliances (on/off) at any time. The resulting vector has n elements, one per appliance in the facility, so that

$$a_i(t) = \begin{cases} 1, & \text{if } i\text{-th appliance is switched on at time } t \\ 0, & \text{if } i\text{-th appliance is switched off at time } t \end{cases} \quad (1.2)$$

for $i = 1, 2, \dots, m$. Assuming that all the appliances in the facility have only two states (on/off), the whole-installation load can be reformulated as:

$$P(t) = \sum_{i=1}^m a_i(t)P_i(t) + e(t) \quad (1.3)$$

where $P_i(t)$ is the switched-on load associated to the i -th appliance and $e(t)$ gathers all the consumptions of the devices excluded from the study and the measurement noise. A first attempt of NILM could be to estimate *when* the appliances are switched on no matter what their actual consumptions $P_i(t)$ are, i.e., to obtain $\mathbf{a}(t)$. Thus, this first energy disaggregation could be formulated as

$$f : P(t) \rightarrow \mathbf{a}(t) \quad (1.4)$$

being f a disaggregation function which maps the whole-installation load into a binary vector of states. If, in addition to *when* an appliance is turned on, the actual consumption of the individual load is also needed, i.e., *how* are the appliance-specific loads, a fine-grained decomposition can be formulated as follows

$$F : P(t) \rightarrow \mathbf{p}(t) \quad (1.5)$$

where the energy disaggregation function F maps the whole-installation load into a continuous vector $\mathbf{p}(t) \in \mathbb{R}^m$ whose elements are the appliance-specific loads $p_i(t)$ at time t . This sequence-to-sequence mapping is more complex than function f , since it strongly depends of the nature of the individual loads under study. Hence, if all the appliances are of the type on/off, obtaining the appliance-specific energy demand waveforms from a model similar to Eq. 1.3 is straightforward, since $p_i(t) = \alpha_i(t)P_i(t)$ where $P_i(t)$ is a known consumption that can be a constant load or a well-known waveform. However, some real appliances have more than two states or even an undefined number of states. Taking this into account, several works [29, 32, 33] have categorized appliances in the four types shown in Figure 1.2.

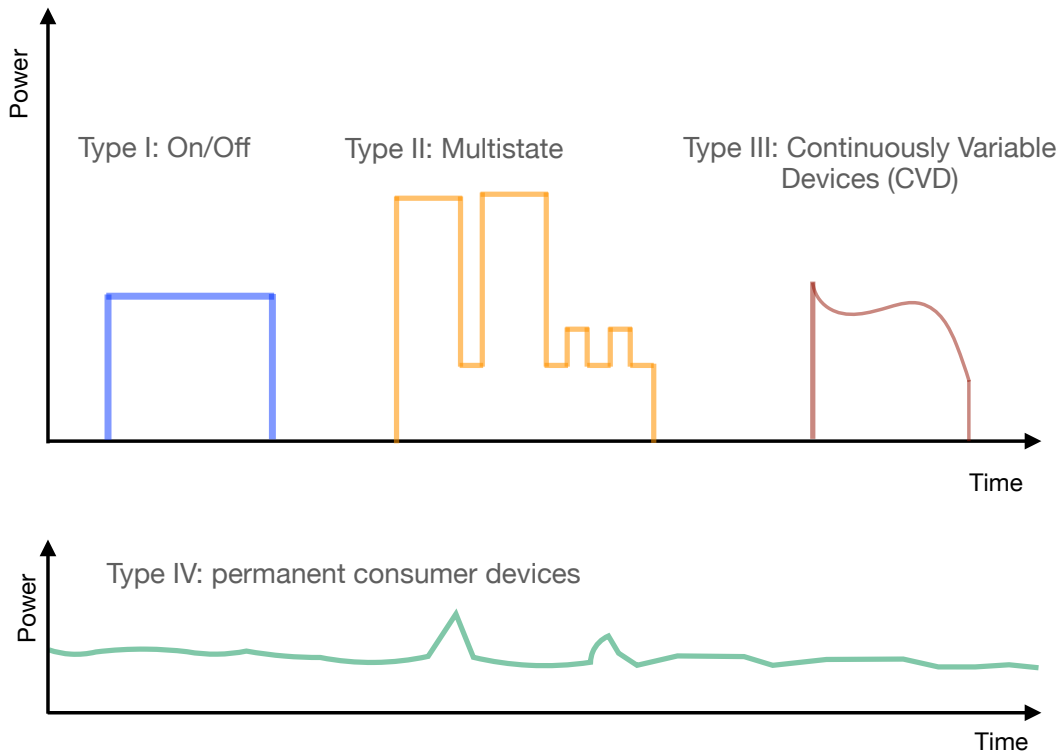


Figure 1.2: Types of appliances in residential facilities.

- *Type-I*. Those devices which have only two operational states (on/off).
- *Type-II*. Devices with more than two states of operation which could be modeled as a *Finite State Machines* (FSM) [29].
- *Type-III*. Appliances with an unknown number of states, also called *Continuously Variable Devices* (CVD). Devices for which users regulate their energy demand in

a continuous way, such as dimmer lights or drills, are included in this category. The human factor makes Type III appliances difficult to identify in the whole-installation load signal.

- *Type-IV permanent consumer devices.* In [33], another type of devices is introduced to the appliance's taxonomy for those devices which remain turned on throughout long periods of time. They named these devices as *permanent consumer devices* since they often draw an approximately constant power without any periodic pattern related to the human activity. Examples of devices under this category are smoke detectors, occupancy sensors or security alarms.

A complete NILM system must be able to deal with all the types of appliances and learn the dynamic behaviors that govern and define the energy demand $P_i(t)$ of each appliance.

The initial formulation of Hart has motivated an intensive research in NILM in recent decades yielding different perspectives to tackle the problem, such as optimization approaches, probabilistic methods, pattern recognition models, and more recently, deep learning techniques that all of them seek to provide users with valuable feedback that improves the overall understanding of the facility. Proper feedback that helps to improve the total efficiency of the facility relies not only on the accuracy of the decomposition, but also on the user's understanding and confidence in the results. The process of interpreting the disaggregation, and therefore of adopting concrete efficiency improvement actions is sped up when users can interact directly with the results and associate them with their previous knowledge about the facility.

In this regard, several works [34–36] have highlighted that NILM systems should satisfy other design principles beyond the accuracy in the load identification in order to maximize the improvement in the efficiency and the user's acceptance of the tool. Among these principles, the requirement of real time NILM where the user can interact with the estimated appliance-specific loads and context data (e.g., ambient temperature, hour, or month.) is one of the most relevant. It is at this point where the theory of visual analytics (VA) converges with the NILM framework, since VA-based techniques perfectly fit on the process of interpreting raw energy decompositions by means of data visualization and interaction mechanisms. The principal VA idea of a *fluid loop*, in which the user tries to solve questions by means of visual representations of the results of an intelligent data

analysis stage and, at the same time, the gained knowledge leads the user towards new hypotheses that he tries to solve again through the visualization, could be applied to the NILM problem. This infinite loop of analysis, graphically represented in Figure 1.3, leads users to more thorough analysis, stronger conclusions and an increment of their trust on the model.

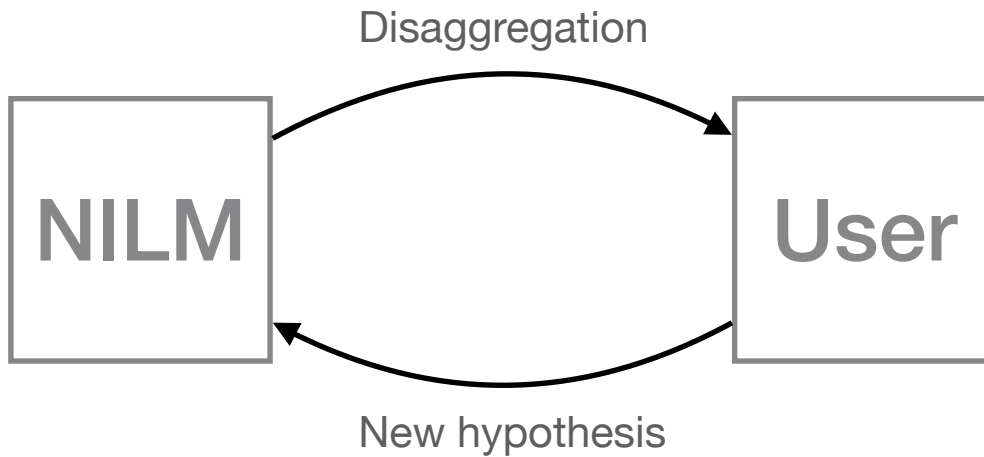


Figure 1.3: Two-way analysis NILM-User.

The aforementioned symbiosis between intelligent data analysis, visualization techniques and interaction mechanisms was formally defined in [3], where a conceptual diagram of VA was proposed. In the case of getting actionable feedback from the NILM outcome, the diagram in [3] can be adopted and modified for the specific load identification task. This is shown in Figure 1.4, where the original diagram has been expanded by integrating new stages that model the NILM process and the facility.

Assuming that boxes represent containers of information and circles represent operations which transform inputs into outputs, the final goal of the VA approach applied to NILM is to increase the user's knowledge about the energy demand of the electrical installation (*energy awareness*), that is represented as K . This increment through the time of the analysis could be defined as:

$$\frac{dK}{dt} = P(I, K) \quad (1.6)$$

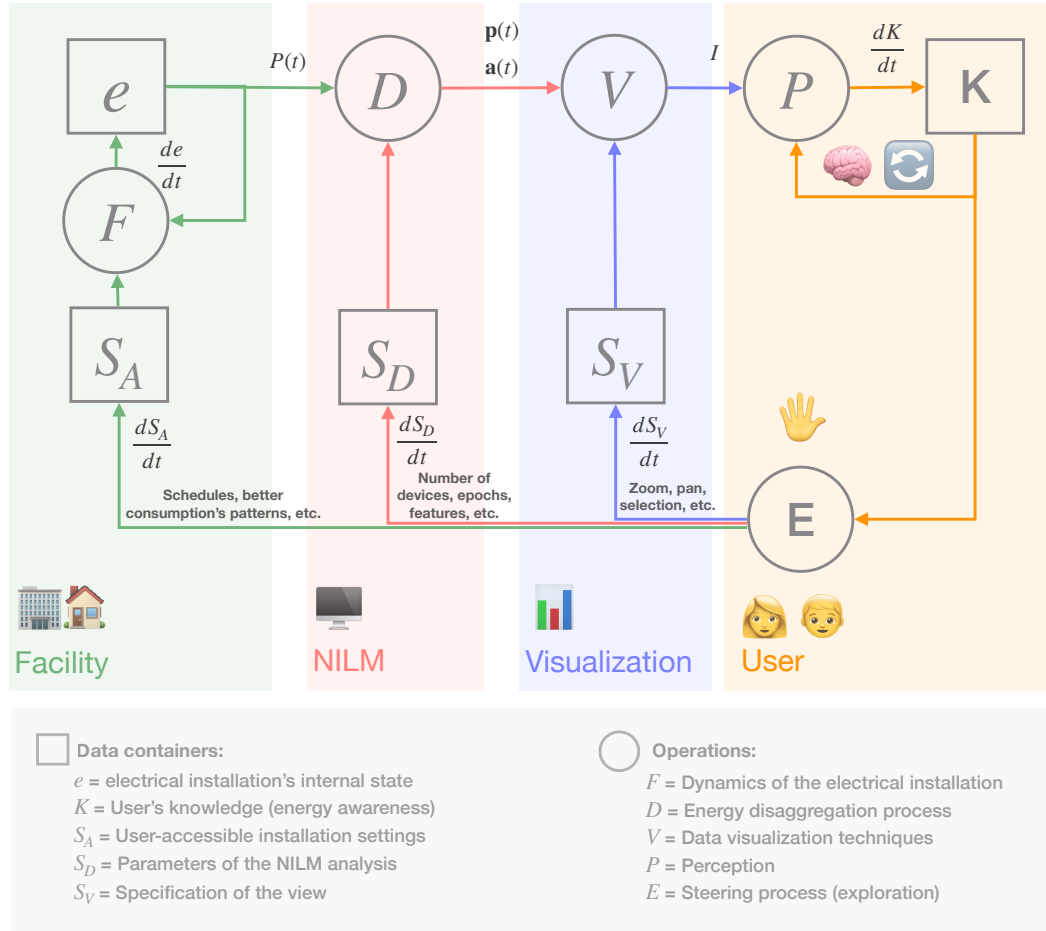


Figure 1.4: Concept model of getting actionable feedback from a NILM system. Inspired by [3]

where the knowledge increment $\frac{dK}{dt}$ is the result of the *perception operation* P of the view shown to the user I and it is also influenced by the user's prior knowledge about the installation K . As the analysis progresses, the initial knowledge about the installation $K(0)$ is increased at a rate of $P(I, K)$. Therefore, the total amount of gained knowledge during the analysis is:

$$K(t) = K(0) + \int_0^t P(I, K) \quad (1.7)$$

The view I is the result of applying data visualization approaches to the obtained load identification and its overall appearance is strongly attached to the specification parameters S_v (e.g. color scales, arrangement of data items in the view, size of elements

and zoom)

$$I = V(\mathbf{p}(t), \mathbf{a}(t), S_v) \quad (1.8)$$

Note that the *visualization operation* V takes the resulting decomposition vectors $\mathbf{a}(t)$ (appliance state identification) and $\mathbf{p}(t)$ (complete estimation of individual consumptions) as inputs. Both decomposition vectors are the outcome of the applied energy disaggregation techniques to the total consumption, that in our diagram are all integrated in the *energy disaggregation operation* D

$$\mathbf{p}(t), \mathbf{a}(t) = D(P(t), S_D) \quad (1.9)$$

where $P(t)$ is the whole-installation load signal and S_D represents the user-driven hyperparameters of the load disaggregation techniques that modulated the resulting decomposition. Hence, the information measured from the facility $P(t)$ is properly transformed by D and V operations, obtaining as a result a view that the user can easily interpret by his powerful perception system P in order to increase his initial knowledge $K(0)$. This gained knowledge may generate new hypotheses and prompt the user to tweak both the disaggregation D and visualization V operations in order to investigate and prove these new questions. The action of steering the data analysis and its visualization is represented in the diagram by means of the action E , whose purpose is to *explore* new hypotheses, modulating the outcomes of the operations D and V through modifying their specifications S_D and S_V , respectively.

$$S_D(t) = S_D(0) + \int_0^t E(K) \quad S_V(t) = S_V(0) + \int_0^t E(K) \quad (1.10)$$

After steering D and V according to the new questions to be answered, a new view I is generated closing the loop. This constant feedback increases more and more the knowledge K , making the user increasingly aware about the energy demand and more expert in the overall behavior of the facility. From this gained awareness, the decision-making process of taking actions that lead the system to a state of better efficiency is significantly enhanced.

All possible actions that can be taken on the electrical installation under study are represented in the container S_A . An action within S_A could be any affordable action to users which implies a reduction in the energy demand. Examples of these actions could be a better design of the schedules of activity in large buildings, finding malfunctions in the downstream devices or an improvement of the user's consumption patterns. In the same way as S_D and S_V , the changes in the actions $\frac{dS_A}{dt}$ are governed by the operation E :

$$S_A(t) = S_A(0) + \int_0^t E(K) \quad (1.11)$$

$S_A(t)$ actions modulate the internal dynamics F which defines the internal state e of the facility

$$e(t) = e(0) + \int_0^t F(e, S_A) \quad (1.12)$$

Thus, the internal state of the system $e(t)$ is driven by the function F which is also modulated by the user. The final goal is to prompt the user to properly modulate F so that the state of the installation is driven to efficient behaviors that reduce the waste of energy [36].

With this in mind, the main goal of the present doctoral thesis will be the development of load identification techniques D susceptible of being integrated into interactive applications where the decomposition is visualized by means of data visualization techniques V and efficient ways of steering both the visualization stage and the decomposition stage by means of their specifications S_V and S_D .

1.3 Objectives

Considering the suggested problem, the main objectives addressed in this doctoral thesis can be summarized in the followings points:

1. Review of the literature on ML and DL techniques, focusing on those approaches related to energy efficiency in processes and large buildings. In this review, data

disaggregation algorithms applied to electrical consumption will be of special interest.

2. Review of the state of the art in data visualization and visual analytics techniques susceptible to be applied to the field of energy efficiency in order to improve the energy awareness of the user.
3. Development and application of visual analytics techniques to the NILM problem in order to create intuitive visual representations for interpreting and finding patterns in consumption data from large facilities, that could be useful for energy planning, general condition monitoring and detection of malfunctions or failures in the system.
4. Development and application of NILM methods based on modern deep neural networks to real electrical consumption data from large installations, emphasizing their possible integration into visual analytics applications and their possible deployment in real facilities (costs, computational requirements, etc.).
5. Exploring the potential functionalities of NILM methods in daily maintenance tasks of large electrical installations, such as redundancy mechanisms that allow the detection and location of failures.
6. Exploring and identifying potential uses of the proposed DNN-based models and VA techniques in other areas of application where the expert knowledge and the interpretability of the results are also crucial.

1.4 Structure of the document

The document is divided into several chapters that group different areas of the research. These chapters are listed below with a brief description of their contents.

In chapter 2, the fundamentals of the NILM framework are reviewed and the state of the art in NILM is presented.

In chapter 3, the differences between residential and non-residential facilities in terms of energy demand are explored, and a novel unsupervised NILM approach based on the *Non-negative Matrix Factorization* (NMF) algorithm, the *OLAP data cube* structure and data visualization techniques is suggested as a visual analytics model for NILM in large buildings. The model is trained and tested using data from a real hospital.

In chapter 4, novel *deep neural networks* (DNN) models are explored and adapted to our visual analytics paradigm and the nature of the large buildings by means of a novel *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE) model, in which *conditioning mechanisms* can be included in order to enable the user to modulate the behavior of the model according to his intentions. The resulting model is applied to real data from residential and non-residential facilities.

Finally, in chapter 5, the final conclusions of the research are summarized and the future work within the possible open lines of research is described.

NILM FRAMEWORK AND DISAGGREGATION TECHNIQUES

In this chapter, the fundamentals of NILM systems design are reviewed. For that purpose, a standard NILM framework will be presented, as well as a review of the most relevant disaggregation techniques since the formulation of NILM in the 1980s, including novel and powerful approaches based on deep neural networks. Although an exhaustive analysis of all the techniques is beyond the scope of this manuscript, the indispensable NILM concepts and approaches related to our contribution are reviewed. Finally, in this chapter, the open issues and challenges of NILM will be also enumerated.

2.1 NILM framework

The process of designing and developing a NILM system, in the same way that other intelligent data analysis processes, entails a set of interconnected tasks until an appropriate outcome is reached, which, in the case of NILM, is an accurate and interpretable decomposition of the total load of a facility. In order to model that process, the well-known CRISP-DM method [4] provides a methodology and standardizes these steps in the workflow shown in Figure 2.1.

In the initial formulation of NILM of Hart [29], as well as in subsequent works [32, 34, 37–39], a methodology similar to CRISP-DM has been adopted to explain the

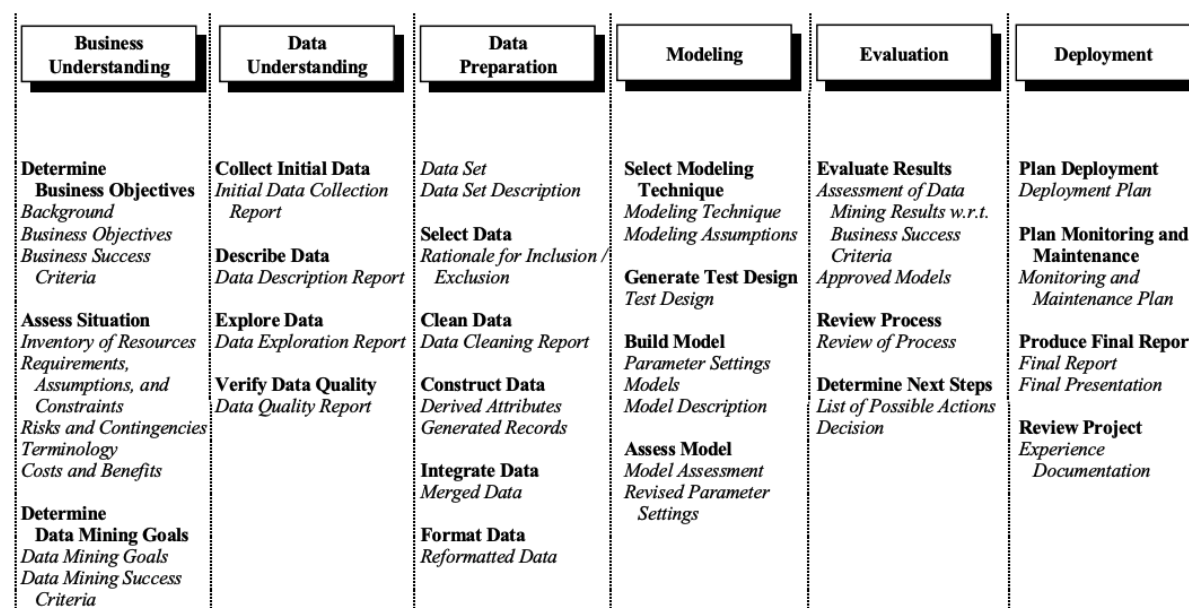


Figure 2.1: CRISP-DM steps defined in the original article [4].

NILM design process, suggesting the framework shown in Figure 2.2. The proposed steps bring together all the processes or tasks involved in NILM, such as the acquisition and preparation of data needed to train energy disaggregation models, the different algorithms for energy disaggregation and their usage to infer the individual consumption loads from the main consumption, as well as all the procedures required to deploy the designed NILM systems in real facilities.

These stages are tightly coupled with each other, and therefore the decisions taken at one stage might affect to the rest of them. As an example, the nature of the information collected in the *data acquisition* step largely determines the following *feature extraction* process and disaggregation techniques used for carrying out the final decomposition. Hence, in the NILM design process, none of the stages can be omitted and all of them should be taken into account.

2.1.1 Data acquisition

The first step in any data analysis problem is to obtain the data needed to resolve the task under study. In the case of NILM, the minimum information needed is the aggregated whole-installation load and, depending on the subsequent algorithms of disaggregation, the signals from the individual consumptions and context information (e.g., states of each appliance, ambient temperature or occupancy) could be also required.

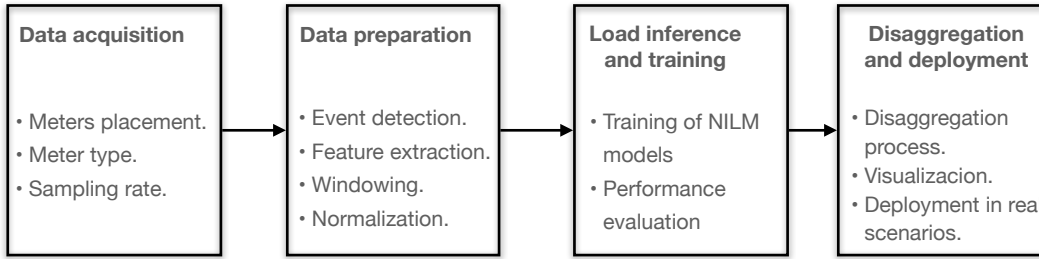


Figure 2.2: Diagram of NILM framework.

The process of measuring consumptions from a facility is not straightforward, since it must ensure the collection of sufficient information on the individual and main loads, and largely determines the cost of implementation, as well as the resolution of the input data for the disaggregation algorithms. At this point, the designer must answer the following questions: *What is the required sample rate? How many meters should be installed? Where should the meters be placed?*. These questions are strongly related to the broad spectrum of available meters [32, 40], which can be divided in:

- *Single function meters*, which are the most widespread metering systems in residential facilities and the most economical option. They usually have only a single channel to measure the active power and their sampling time resolution ranges from 1 s to 60 s.
- *Multifunction meters* which not only measure the active power, but also integrate several channels for other electrical variables, such as currents, voltages, or even harmonics and harmonic-based metrics, as the *total harmonic distortion (THD)* used to monitor the quality of the energy supply. These multifunction energy meters are typically installed in industrial environments or large buildings, where other aspects of the facilities, apart from the energy consumption, play an important role.
- *High-frequency meters*. They can be both multifunction and single function, their sample rate is above 1 MHz, and they are the most expensive ones, since they are typically built with custom hardware. These types of high-frequency meters can sample fine-grained waveforms with sufficient resolution to detect transient states and electric noise caused by changes in the state of the downstream appliances or devices.

Depending on the type and number of the installed meters, the nature of the collected information may vary. Thus, high-frequency multifunction meters enable the development of algorithms based on transient features or harmonics, while if the meter installed is a single function device, subsequent energy disaggregation techniques must be able to extract the individual consumptions from low resolution input data, which makes the disaggregation more difficult.

In the design of a data acquisition module for NILM, we must distinguish between residential facilities, where only a few devices are turned on at same time, and large non-residential electrical installations, such as public buildings, where the whole-installation consumption aggregates, simultaneously and in a continuous fashion, a wide range of devices and circuits of different natures. In the case of residential facilities, load disaggregation could be addressed by means of single function meters as it is shown in [8, 41, 42].

Regarding data acquisition systems in large buildings, they are often the result of a trade-off between: a) installing only a single meter at the supply connection and dealing with a more difficult disaggregation problem; or b) assuming the set-up costs involved in the installation of several meters along the electrical network for reducing the complexity of the subsequent load identifications. In the case of large electrical installations, high-frequency meters are highly desirable, since they are able to measure complete waveforms for identifying anomaly events in the electrical networks caused by malfunctions or demanding start-ups of industrial subsystems that involve, for example, powerful electrical motors.

2.1.2 Data preparation and feature extraction

Before processing the data through energy disaggregation algorithms, NILM systems integrate a previous stage where the data are preprocessed and, in many cases, relevant features are computed from the raw energy consumption signals. In early stages of NILM, disaggregation algorithms are focused on event-based approaches, illustrated in Figure 2.3, where the energy disaggregation relies on the following two-step preprocessing: 1) identifying events from the whole-installation energy demand waveform; and 2) characterizing each detected event by means of a vector of features. Once the events are detected, they are associated with the appliance that triggered them, using energy disaggregation algorithms that take as input a representative set of features computed

on them.

Event detection

Event detection algorithms track the signatures in the aggregated electrical variables looking for the time instants $e^{(i)}$ where an appliance changes its state (e.g., a rise in the main energy load caused by *off* to *on* appliance event). The event detection process, illustrated in Figure 2.3, extracts from the whole-installation consumption sequence $P(t)$, assuming it is measured at $t = 1, 2, \dots, n$, a set of time instants tagged as events $\{e^{(1)}, e^{(2)}, \dots, e^{(i)}\}$, where $e^{(i)} \in \{1, 2, \dots, n\}$.

A naive event detector would be to compare the power level between two contiguous samples from the aggregated consumption, and tag the sample as an event when the difference exceeds a user-defined threshold. Changes in the state of large consuming appliances produce sharp jumps in power levels, but the signatures of devices with lower energy demand may be identical among them, or may even be confused with electrical noise. For this reason, probabilistic models in [43] are more convenient event detectors when low-power appliances are involved. Other approaches are able to detect the events in the frequency domain by monitoring changes in harmonics energy computed from the whole-installation power signal over the time [44–46].

Feature extraction

In the time instants $e^{(i)}$ where events take place, the signal bears the most relevant information about the states of the appliances. The information around events is typically synthesized by means of descriptors, which can be divided into steady-state or transient features [32, 47]:

Steady-state features. Steady-state analyses calculate a vector of features $\{F_1^{(i)}, F_1^{(i)}, \dots, F_m^{(i)}\}$ per event, using only the sequence $P(t)$ in the interval $[e^{(i)}, e^{(i+1)}]$ between two contiguous events. During this sequence, all appliances remain in the same state (see Figure 2.3). Typically, power levels in real power $P(e^{(i)})$ and reactive power $Q(e^{(i)})$ are usually computed as steady-state features $F_m^{(i)}$, assuming that these values are constant in the interval $[e^{(i)}, e^{(i+1)}]$. In [46, 48–51], only the active power was used to characterize the steady-states, but this method fails to detect appliances with similar consumptions. In order to tackle this shortcoming, a wider spectrum of electrical features $F_m^{(i)}$ per event has been also proposed as steady-state features:

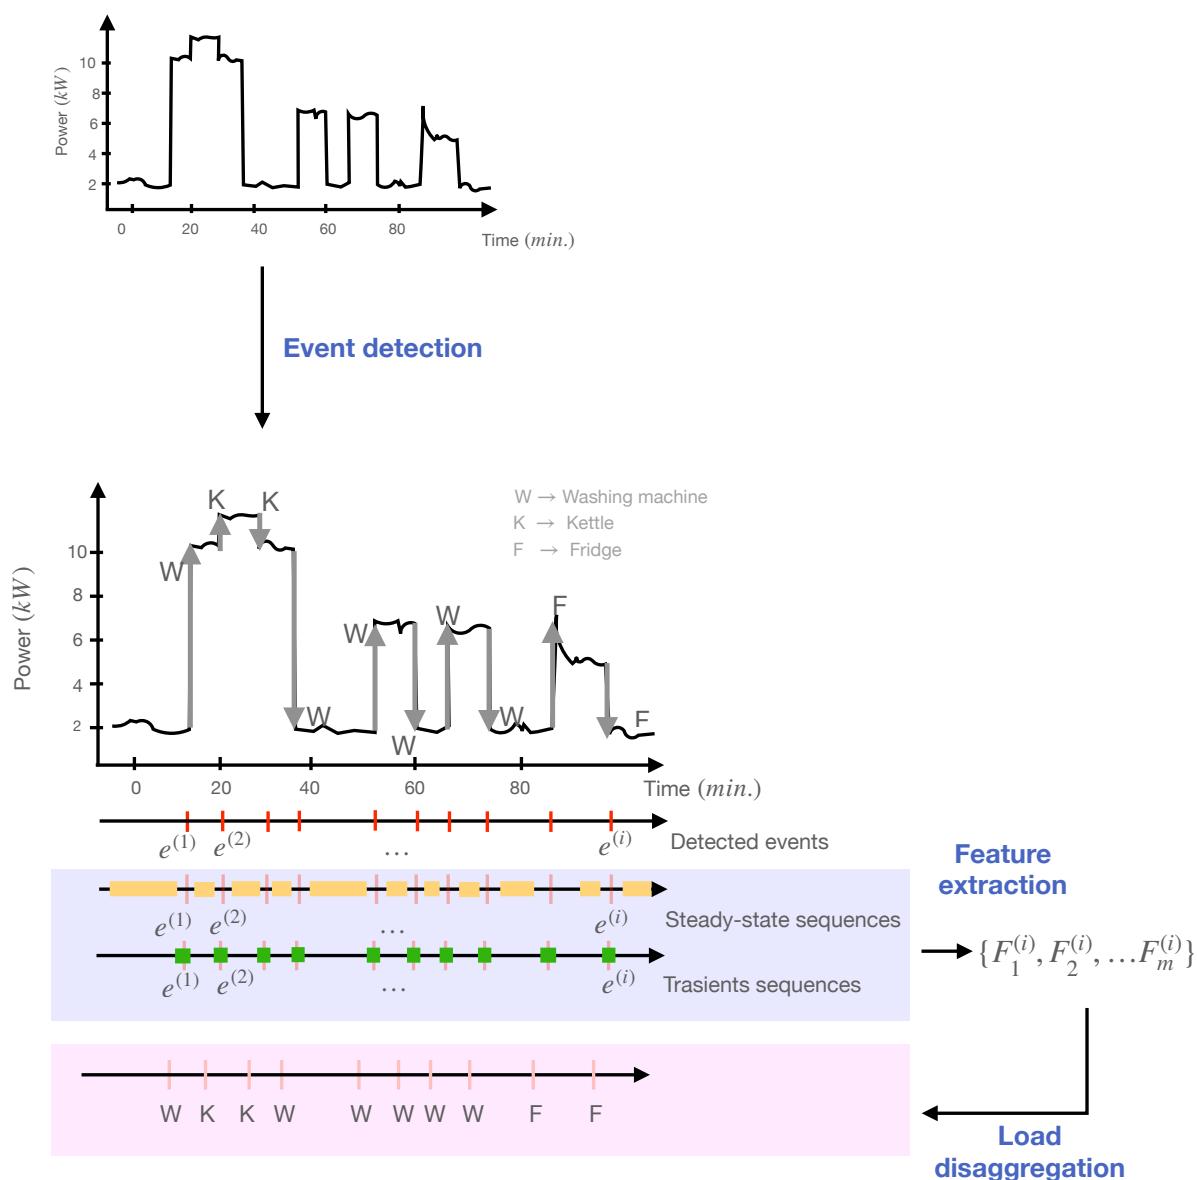


Figure 2.3: Event-based energy disaggregation process. After the event detection, each event $e^{(i)}$ is identified by means of features computed on the *steady state* or *transient* sequences.

- Simultaneous changes of $P(e^{(i)})$ and $Q(e^{(i)})$ power. This leads to the bidimensional set of features $\{P(e^{(i)}), Q(e^{(i)})\}$, known as $P - Q$ plane [29].
- Peak value of voltage $V(t)$ and current $I(t)$ for $t \in [e^{(i)}, e^{(i+1)}]$ [52, 53].
- *Root Mean Square Error* (RMSE) of the voltages $V(t)$, currents $I(t)$ and power factor for $t \in [e^{(i)}, e^{(i+1)}]$ [52, 53].

- $V - I$ raw trajectories [54, 55] made of the complete sequence $V(t)$ and $I(t)$ for $t \in [e^{(i)}, e^{(i+1)}]$.
- Frequency domain features, such as harmonics [56, 57] or *Total Harmonic Distortion* (THD) [58] computed on sequence $P(t)$ for $t \in [e^{(i)}, e^{(i+1)}]$.

Transient features. Transient-based feature extractors assume that the sequences $P(t)$ in the interval $[e^{(i)} + c, e^{(i+1)} - c]$ around the detected events (green segments in Figure 2.3) contain enough information to accurately describe the events, without taking into account the sequences of signal between transients. Hence, in the same way as in steady-state analyses, a vector of features $\{F_1^{(i)}, F_1^{(i)}, \dots, F_m^{(i)}\}$ is computed over the local neighborhood of each event.

Basic descriptors, such as the slope, duration and peaks of current or power during the transient samples, were initially defined in [46, 59, 60]. The major drawback of these initial transient features is that they strongly depend on the appliance type and the context, which significantly reduces their scalability [32].

Spectral envelopes based on the coefficients of several harmonics of the *Short-Time Fourier Transform* (STFT) proposed in [44] have also shown to be powerful and discriminative transient features. The bandwidth of the spectral envelopes is limited by the window size on which the frequency bands are computed, so that monitoring high and low frequency signatures at the same time implies large windows, increasing the computational cost of the extraction, and therefore decreasing its flexibility for real time NILM applications and their deployment in real scenarios [47].

In this regard, the *wavelet-based* analysis provides a more flexible feature extraction, since it shows a better frequency/time trade off and reduces the computational requirements [61, 62]. Despite STFT and wavelet-based transformations have proven to be useful as appliance-specific features, both methods require input measurements with high sampling rates [60], and therefore expensive data acquisition modules, which is contrary to the current trend of low-cost approaches.

In the last decade, modern NILM systems have increasingly replaced the extraction of user-defined features based on event detection by novel *deep neural networks* (DNN) architectures, which automatically extract the features needed for the subsequent load identification from the raw whole-installation electrical variables [42, 63–65]. Instead of

using features specifically designed for certain appliance behaviors, DNN-based NILM models get rid of any assumptions and learn convenient features by stacking layers that increasingly identify complex and meaningful patterns, resulting in a compositional hierarchy of features known as *representational learning* [66].

Since the user-defined feature extraction in DNN-based approaches is removed, the data preparation stage is simplified to those tasks needed to adapt the energy consumption measurements into a convenient format to train the networks and to infer the individual consumptions, when these models are deployed. An example of these operations is the *windowing operation* [63] shown in Figure 2.4, which transforms the measured energy consumption times series into an appropriate input format for the DNNs, breaking them down into sequences of the same length and normalizing them.

2.1.3 Load inference and learning

Once the data is prepared, the features or windows of the aggregated signal are processed by load disaggregation algorithms in order to obtain the states and energy demand of the individual downstream consumptions. Conceptually, load disaggregation models were defined in section 1.2 as f , when they only estimate the state of the appliances at any time $a_i(t)$, and as F , if they estimate the complete power waveform of the individual loads $p_i(t)$. The different models used to implement functions F and f can be divided into *supervised techniques* and *unsupervised techniques*.

Supervised approaches require labeled data, which in NILM means that data from the individual appliances must be available. By contrast, unsupervised approaches learn how to extract the individual appliances only from the main energy demand. In Figure 2.5, a detailed taxonomy of the disaggregation techniques found in NILM literature is graphically presented. This taxonomy will be thoroughly reviewed in subsequent sections of this chapter.

Finally, this stage of the development of NILM systems also includes the performance evaluation of the trained disaggregation approaches by means of metrics specifically defined for NILM, which will be reviewed in section 2.4.

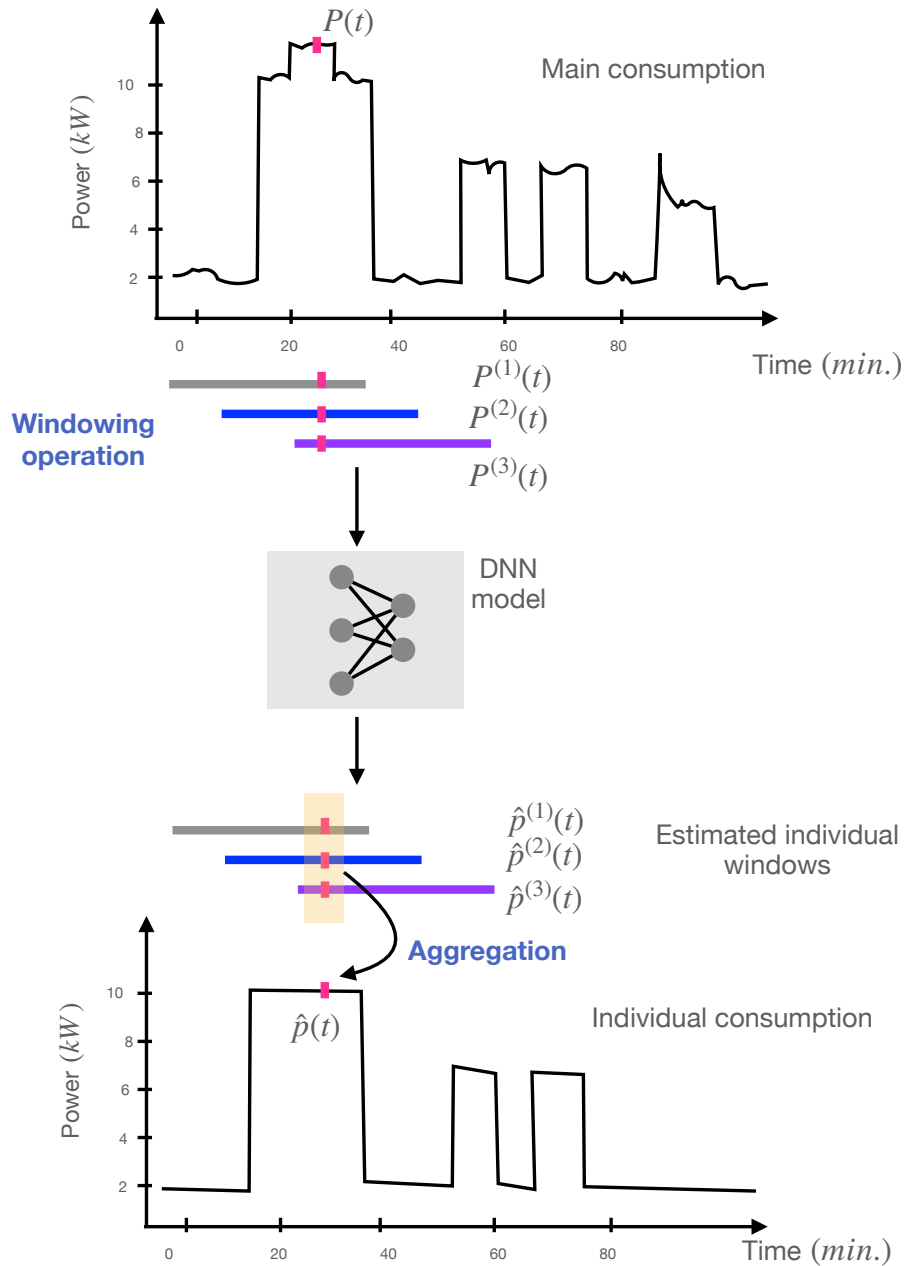


Figure 2.4: DNN-based energy disaggregation process. The main consumption of the facility is split in a set of overlapped sequences, using a windowing operation. The resulting windows are processed by a trained DNN model to produce a set of overlapped windows of the estimated individual consumption. The final individual consumption waveform is computed by aggregating out windows of DNN model by means of an aggregation function, such the mean.

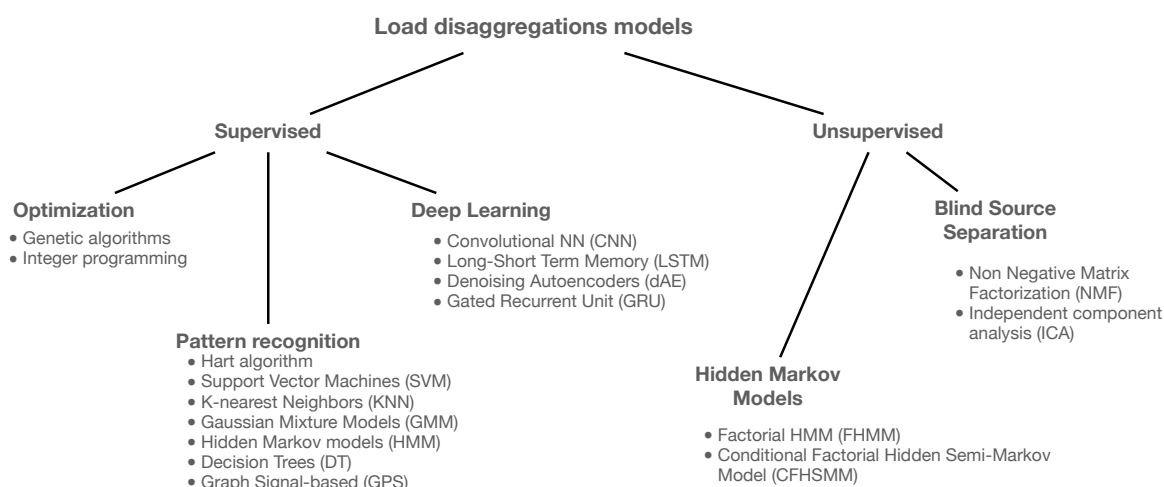


Figure 2.5: Taxonomy of load disaggregation methods.

2.1.4 Disaggregation and deployment

All the operations after the load identification belong to this category, and they are typically in charge of: 1) completing the disaggregation, 2) the deployment in real facilities and 3) providing the user with a proper information about the amount of energy demand per device. Some load disaggregation algorithms need a post-processing phase, where the load identification is transformed in order to obtain a final and interpretable disaggregation. For instance, in the case of some DNN-based methods, several estimations of an individual consumption for the same instant are made, because their input consists of overlapping windows. Thus, an aggregation operation should be applied to these estimations to obtain a disaggregation of the same length as the total consumption.

The development environment (e.g., data or computer equipments) usually differs from the real environment where the NILM system will be finally deployed, so that NILM models have to be prepared to be transferred into the real scenario. Examples of tasks involved in this adaptation could be the integration of NILM models in commercial energy management systems or the creation of a web service to which smart meters could send the aggregated signals and, after being processed, receive back the estimated individual consumptions.

This stage also includes all tasks focused on improving user acceptance of the decomposition. In [36], authors highlighted that the energy disaggregation by itself is not a compelling application, since it does not lead directly to actionable feedback capable of

improving the energy efficiency. Therefore, the authors suggested an “application-centric” paradigm, where the research on NILM should also emphasize the design of novel and powerful applications based on NILM, rather than focusing only on improving accuracy. This is perfectly aligned with our problem formulation suggested in section 1.2, where the need of a VA mechanism for improving the energy awareness of the user was proposed. Thus, all the tasks that improve the interpretation of the resulting energy decomposition, such as its integration into a VA paradigm, could be included in this stage.

2.2 Supervised disaggregation algorithms

According to our problem formulation, supervised NILM algorithms can be defined as functions F and f capable of mapping the extracted features after the event detection, or directly the windows of the aggregated signal $P(t)$ to an estimation of the appliance states $\mathbf{a}(t)$ (in the case of functions f) or the appliance consumption waveforms $\mathbf{p}(t)$ (in the case of F models). These supervised mapping functions are the result of a training process, where information from the individual appliances (labeled data) is required. The optimal functions obtained after the training stage result from the optimization of a *loss function* of the form

$$\mathcal{L}(p(t), \hat{p}(t))$$

which typically compares the ground truth appliance consumption waveform $p(t)$ and their corresponding model estimation $\hat{p}(t)$. The branch of supervised approaches in the taxonomy shown in Figure 2.5 will be further explained in the following sections, paying special attention to optimization, pattern recognition and deep learning models.

2.2.1 Optimization methods

Optimization methods address the load identification as a *Combinatorial Optimization* (CO) problem. Considering $\hat{\mathbf{y}}$ the feature vector computed from a detected event or state between events, optimization methods search the best matching vector of features \mathbf{y}_i within an available appliance features database. This optimization can be formally expressed as:

$$\operatorname{argmin}_i \|\hat{\mathbf{y}} - \mathbf{y}_i\| \quad (2.1)$$

where i is the optimal index of the appliance database and it identifies the appliance state to switch on within the output states vector $a_i(t)$. When only one appliance is switched-on, finding the index i is straightforward, since it is a one-to-one matching [32]. Nevertheless, when the number of appliances that can be activated at the same time and the number of states per appliance increase, the aforementioned optimization becomes complex, since all possible combinations between appliances and states have to be considered.

Therefore, researchers in [52, 67–70] have tackled the CO optimization problem using more complex optimization algorithms, such as *integer programming* or *genetic algorithms*. Several authors [29, 32, 71] have pointed out that these CO approaches present the following problems:

- The optimization problem presented in Eq. 2.1 becomes computationally unfeasible when the number of appliances and the internal states increase.
- Generating an appliance-specific profile for all the possible contexts (e.g., localizations, type of devices or operations modes) is quite difficult and it significantly restricts the scalability of CO problems.
- Some appliances present similar signatures in $P(t)$ making the optimization difficult.
- These models only estimate the states of each of the appliances $\mathbf{a}(t)$, but final users are not only worried about when the appliances are turned on, but also about how much energy they have spent.
- The optimization becomes more difficult when appliances out of the database are plugged-in.

2.2.2 Pattern recognition methods

In contrast to CO models, which search the best matching profile in a database, *pattern recognition* NILM methods are based on continuous functions that map regions of the feature space to possible states of the appliances. An intuition of the pattern recognition

framework is shown in Figure 2.6, where a function f learns the boundaries of the feature space, attaching them to an appliance-specific state.

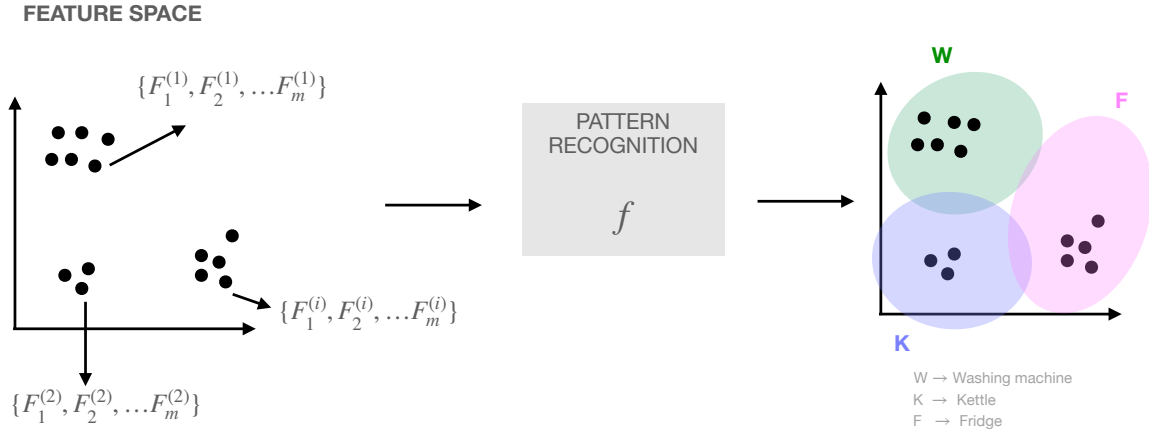


Figure 2.6: Conceptual diagram of pattern recognition NILM models.

Hart defines a first pattern recognition approach in [29], where a feature space formed by the P - Q plane is delimited into regions by means of a clustering technique, in which each cluster is assigned manually or automatically to a specific appliance. Projecting an unknown event vector of features from an aggregated sequence into the P - Q plane and computing the best matching appliance cluster is how the loads are identified in the inference phase.

The major limitations of Hart’s approach are the occlusions on P - Q plane of the less-consuming appliances and the incompatibility with continuously active appliances [71]. Subsequent works have further developed pattern recognition methods, integrating other ML techniques capable of mapping event’s features to appliance states. Among them, in [72], a naïve Bayes classifier tags the event’s features. In [73–76], the recognition of the event features is carried out by more complex ML models, such as *Support Vector Classifiers* (SVC). In addition, *Decision Trees* (DT) [77], *K-Nearest Neighbors* (KNN) [78] and shallow *Neural Networks* (ANN) [53] are also applied to attach feature vectors to specific load appliances.

Contrary to the aforementioned purely event-based methods, several authors started taking into account the *dynamics* of the individual consumptions, using the previous state $\mathbf{a}(t-1)$ in combination to the event’s features to classify an incoming event [79]. Hence, probabilistic methods such as Hidden Markov Models (HMM) [80] emerged, which consider the signature of each appliance as the observation of an internal vector

state $\mathbf{a}(t)$, which is estimated according to the previous state $\mathbf{a}(t - 1)$ (Markov property). Recently, researchers in [41, 81] have suggested embedding the aggregated signal $P(t)$ into a graph by means of the graph-based signal processing (GSP) technique, resulting in an effective NILM model for low-resolution scenarios.

2.2.3 Deep learning methods

The performance of the aforementioned methods relies on a user-defined set of features computed on events extracted from the whole-installation consumption waveform, assuming that they are correct or appropriate to the load identification purpose and robust against changes in the context (e.g., changes in facility, new appliances or seasonality). In [33], authors concluded that extracting a reliable vector of features capable of representing all types of appliances is not straightforward, at least with state-of-art combinations of steady-state and transient features, so that a more efficient and flexible representation of the whole-installation consumption should be further explored.

At this regard, *Deep Learning techniques* (DL) have shown an extraordinary capacity to automatically extract discriminative features in image classification tasks [17], natural language processing [19, 82] or time series modeling [83]. The idea of allowing DL-based models to automatically compute features from the aggregated signal and estimate directly the states of individual appliances and/or their energy demand was translated to NILM for the first time in [64, 71]. Authors showed the feasibility of DL-based NILM models by adapting several modern neural network architectures to the NILM problem, such as *Convolutional Neural Networks* (CNN) [84], *Denoising Auto-Encoders* (dAE) [85] and *Recurrent Neural Networks* (RNN), like *Long Short-Term Memory* (LSTM) models [86, 87]. Proposed DL-based NILM models in [64, 71] outperformed previous state-of-art approaches based on HMM-based methods, which were a milestone in NILM literature, and encouraged other authors to further develop the proposed DL-based architectures:

- *CNN* approaches rely on the hypothesis that, by stacking convolutional layers, the resulting architectures are sufficiently powerful in finding relevant patterns or features to identify appliance-specific signatures in the aggregated signal. In [88], authors pointed out that a deep CNN architecture with skip connections also called *residual CNN* [89] is capable of accurately estimating when the appliances are turned on (vector $\mathbf{a}(t)$). In [90], the same residual CNN architecture is also suggested, but in this case the resulting model is a sequence-to-sequence transfor-

mation, which estimates the complete energy demand of each appliance. Recent works [91, 92] have shown that adding extra information apart from the active power, such as other electrical variables or contextual information to the input improves the accuracy of the energy disaggregation.

- *LSTM* approaches are state-based models similar to HMM-based models, but they provide more capacity for modeling time series, since they integrate a flexible continuous internal state rather than discrete hidden states. Moreover, the internal state of LSTMs depends not only on the previous state, but also on the previous history of states, avoiding the Markov assumption. The study of LSTM-based NILM models is an active research topic, where the major contributions are aimed at: integrating LSTM cells into a dAE architecture [93]; exploring novel bidirectional LSTM (biLSTM) layers [94] which take into account the past samples to estimate the internal state, but also the future samples; or reducing the complexity of the architecture using *Gated Recurrent Units* (GRU) [95].
- *dAE* approaches, by contrast, tackle the energy disaggregation as a denoising problem where the whole-installation load is the noisy signal to be “cleaned” in order to obtain the appliance loads. Several works [42, 64, 96–98] have reported good results in processing overlapping segments or windows of the total consumption using filter functions based on dAEs. These dAE-based approaches are *sequence-to-sequence* models, in other words, they extract an output signal with the same length as the input signal. The sliding window approach in dAE-based models is prone to be sensitive to the input window size and the percent of overlapping between windows. In [65], a sequence-to-point architecture was presented showing invariance to the selected input window size and overlapping ratio when the output point is placed in the middle of the input window.

2.3 Unsupervised disaggregation algorithms

Previously mentioned approaches make use of some sort of *a priori* information (appliance data or information about the installation, such as the number of appliances or appliances’ features) in their training process before deployment. Since the first steps of NILM, one of the NILM challenges has been to reduce the set-up costs as much as possible. For this reason, the development of completely unsupervised models capable of providing a load monitoring without expensive modifications in the installation and a

laborious training process is the final goal of the NILM community. In the following section, we will review the most relevant advances in unsupervised disaggregation methods, following the taxonomy shown in Figure 2.5.

2.3.1 Unsupervised HMM-based models

The initial probabilistic HMM approach for NILM [80] was expanded to a factorial variant in [8] called *Factorial Hidden Markov Models* (FHMM), where the internal state is no longer a unique value, but instead a vector of states corresponding to all states of all appliances. Context information, such as the ON-duration time distributions and information about the time and day of week, is also included to the model, so that two new perspectives of FHMM are also suggested: 1) *Conditional Factorial Hidden Markov Models* (CFHMM), in which the internal state not only depends on the previous one, but also on the time and day of week; and 2) *Factorial Hidden Semi-Markov Model* (FHSMM), in which a gamma distribution over the hidden ON-duration states is imposed. These conditional approaches outperform vanilla FHMM when the number of appliances is increased.

The major drawback of FHMM is that the inference of the internal states is prone to local minima. Subsequent works have further improved the process of internal states inference by means of a convex inference technique called *Additive Factorial Approximate Maximum A Posteriori* (AFMAP) [99]. Although FHMM-based methods achieve high accuracies for appliances of types I and II, they cannot handle continuously active appliances, since a finite discrete set of states per appliance is assumed. In addition to this, when the number of appliances and states per appliance grows, the complexity of FHMM-based is also increased, causing a drop in performance [37].

2.3.2 Source Separation methods

Source Separation methods tackle the NILM problem as a single-channel *Blind Source Separation* task (BSS) [100], where the aggregated signal is considered a mixture of unknown (blind) sources which, in the case of NILM, are the appliance-specific loads. BSS NILM methods were introduced in [101] using a *discriminative sparse coding* to estimate the sources. Authors introduced into the decomposition some energy demand priors to reduce the complexity of the source optimization, such as positive restrictions due to the nature of energy demand data. The non-negative constraints were further

studied in [102], where a tensor variant of the *non-negative matrix factorization* (NMF) [5] algorithm is suggested to provide better representations of the dependences between appliances.

2.3.3 Unsupervised graph-based signal processing

In [103, 104], authors explored the feasibility of unsupervised GSP approaches based on signal clustering and a pattern matching. The requirement of manually tagging each estimated appliance load and the loss of accuracy when several appliances draw similar amounts of power are some of the shortcomings of the unsupervised GSP-NILM models.

2.4 Evaluation of NILM methods

Assessing the performance of load disaggregation models is one of the open issues on the NILM field, due to the lack of agreement in the evaluation procedure. Most of the published works evaluate their suggested models with their own data and using different metrics, hindering their reproducibility on different scenarios. Although some authors [105, 106] have made steps towards the reproducibility of NILM models by developing a widely accepted toolkit which integrates the most used datasets and metrics, a standard procedure of evaluation is still to be defined.

Regarding the metrics used in the literature to evaluate the performance of NILM methods, they are commonly divided into: 1) classification metrics, which measure the accuracy of the NILM methods in determining appliances states $\mathbf{a}(t)$; and 2) energy estimation metrics, which measure how well the waveform of each individual consumption $\mathbf{p}(t)$ is estimated by the energy disaggregation algorithms.

Classification metrics

The first steps on NILM research were focused on event-based approaches and they were evaluated by event-centric metrics. For instance, Hart, in [107], suggested the fraction of the correctly classified events and the percent of the total energy located by the correctly classified events as event-based metrics. With the emergence of the *eventless* NILM techniques, such as HMM-based or DNN-based NILM models, a new paradigm of binary classification appeared, where the estimated individual loads are evaluated on a sample-by-sample basis, rather than by correctly classified event. *Eventless* approaches can be divided into those models which produce a binary output, such as HMM-based

or DNN-based trained to learn the binary states vector $\mathbf{a}(t)$, or those which estimate energy demand of the appliances, such as dAE-based approaches. The later ones could be seen as binary classifiers if a threshold operation is applied on their continuous output estimations.

The resulting binary states vector can be evaluated by means of well-known classification metrics used in other classification problems. Common metrics used in NILM literature are summarized in Table 2.1. Since the NILM problem can be considered as a *multiclass* problem (being the on/off state of each appliance a class), each of the metrics of Table 2.1 could be applied on each class, if we assume the paradigm *one vs all*, and then the resulting values per metric could be averaged in order to obtain an overall performance of the system. This approach is denominated *macroaveraging* [108], and it is complementary to the *microaveraging* approach, where TP, TN, FP and FN values are computed from the entire set of classes, before computing the rest of the metrics shown in Table 2.1.

Metric	Description	Equation
TP	True positives [43, 71, 88, 91, 93, 97, 109]. The number of time slices or samples that the disaggregation techniques have classified as ON or true and they are actually true in the measured individual load. In other words, the number of ON samples correctly identified by the NILM system under evaluation.	-
TN	True negatives [43, 71, 88, 91, 93, 97, 109]. The number of time slices or samples that the disaggregation techniques have classified as OFF or False and they are actually false in the measured individual load. In other words, the number of OFF samples correctly identified by the NILM system under evaluation.	-
FP	False positives [43, 71, 88, 91, 93, 97, 109]. The number of time slices or samples that the disaggregation techniques have classified as ON or True and they are actually OFF in the measured signal. In other words, the number of OFF samples incorrectly classified by the NILM system under evaluation.	-
FN	False negatives [43, 71, 88, 91, 93, 97, 109]. The number of time slices or samples that the disaggregation techniques have classified as OFF or False and they are actually ON in the measured signal. In other words, the number of ON samples incorrectly classified by the NILM system under evaluation.	-
R	Recall [71, 88, 91, 93, 97, 109]. It is also called sensibility and it could be considered as the probability that an ON sample of an individual consumption is correctly classified by the NILM system.	$\frac{TP}{TP+FN}$
P	Precision [71, 88, 91, 93, 97, 109]. It is the proportion between ground truth ON samples in the individual consumption and the total number of samples tagged as ON by the NILM system.	$\frac{TP}{TP+FP}$
F1	F1-score [8, 71, 88, 91, 93, 97, 109]. It is the harmonic mean between P and R and it could be seen a trade-off value between these two metrics. It is recommended for unbalanced individual consumptions (More OFF samples than ON or vice versa.)	$2 \times \frac{P \times R}{P+R}$

Table 2.1: Common binary classification metrics for NILM.

Energy estimation metrics

Energy estimation metrics assess the deviation between the estimated individual load waveforms and the ground truth energy demand of the individual loads. In Table 2.2, the common energy estimation metrics used in NILM research are shown. As in other regression problems from different domains, common energy estimation metrics in NILM are the *Root Mean Squared Error* (RMSE) or the *Mean Absolute Error* (MAE). Although RMSE and MAE are widespread methods in the NILM literature, some authors [37, 97, 105] suggested that if the RMSE is used as the evaluation metric, the comparison between individual disaggregations is not coherent, since the total energy of each consumer has to be taken into account for a fair comparison between low-power and high-power individual loads.

Thus, metrics normalized by the total power consumed of the load under study are commonly adopted, such as the *Normalized Error in assigned Power* (NEP), which is computed dividing the sum of the absolute values of the difference between estimated power $\hat{p}_i(t)$ by the total amount of real power $\sum_{t=1}^T p_i(t)$ consumed by the individual consumption i .

Note that aforementioned energy estimation metrics are computed by appliance, so that they only evaluate one of the appliance-specific estimations at a time. In order to evaluate the overall energy estimation, authors in [71] defined the *Relative Error in Total Energy* (RETE) as an appropriate metric for reporting how accurate is a NILM system in assigning the total amount of energy consumed in a facility.

As it was mentioned, the evaluation of the NILM outcome is a matter of discussion which has led the authors to propose their own metrics or to combine several metrics in order to provide a reliable assessment of the results. The amount of different metrics suggested in NILM literature is huge and an in-depth analysis of all of them is beyond the scope of the present dissertation, so that the metrics presented in Table 2.2 are the most relevant ones and those which we will be applied latter as our evaluation metrics. For a thorough review about NILM evaluation, please refer to [108, 110].

Metric	Description	Equation
T	indicates the total number of time steps of the time sequence to be analyzed.	-
M	denotes the total number of appliances in the study.	-
$p_i(t)$	denotes the ground truth energy demand of i -th appliance at time instant t .	-
$\hat{p}_i(t)$	denotes the estimated energy demand of i -th appliance by the energy disaggregation techniques at time instant t .	-
E	denotes the total energy of the facility gathered by the meter during a time sequence of T time steps.	$\sum_{t=1}^T \sum_{i=1}^M p_i(t)$
\hat{E}	indicates the estimated total energy of the facility by the energy disaggregation techniques during a time sequence of T time steps.	$\sum_{t=1}^T \sum_{i=1}^M \hat{p}_i(t)$
RMSE	Root Mean Squared Error [73, 94, 102, 111] is defined as the standard deviation of the errors in each estimated appliance-specific loads. Note that this metric is computed per appliance.	$\sqrt{\frac{\sum_{t=1}^T (\hat{p}_i(t) - p_i(t))^2}{T}}$
MAE	Mean Absolute Error [65, 71, 90, 91, 96] is the averaged of the absolute difference between the ground truth individual energy demand and its estimation. Note that this metric is computed per appliance	$\frac{\sum_{t=1}^T \hat{p}_i(t) - p_i(t) }{T}$
NEP	Normalized Error in assigned Power [70, 71, 94, 97, 111, 112] is defined as the absolute difference between the ground truth individual energy demand and its estimation normalized by the total energy consumed by the appliance. Other authors used the terminology <i>Signal Aggregate Error</i> [65, 91]. Note that this metric is computed per appliance.	$\frac{\sum_{t=1}^T \hat{p}_i(t) - p_i(t) }{\sum_{t=1}^T p_i(t)}$
RETE	Relative Error in Total Energy [71] indicates the absolute difference between the total measured energy and its estimation, normalized by the maximum between them.	$\frac{ \hat{E} - E }{\max(E, \hat{E})}$

Table 2.2: Common energy estimations metrics for NILM. Note that the first rows of the table indicate the notation used in the metrics.

2.5 Challenges and open issues

Despite the intense research on NILM along the last decades and the remarkable accuracy that some of the aforementioned approaches have achieved, NILM technologies have not had neither the expected impact on society, nor a massive deployment in residential electrical installations yet, even though the use of smart meters are widespread in developed countries. The principal issue of NILM models is that they are difficult to reproduce in environments other than the training benchmark datasets [34, 113]. The reasons of the lack of adoption of NILM systems have been summarized in the following points:

- *Low cost.* Set-up costs of NILM systems are the principal bottleneck of their massive deployment as we have mentioned before. The overall cost of a NILM application highly depends on its data acquisition systems and the type of features used as input of its energy disaggregation technique. Thus, if the NILM system relies on the signals coming from only one meter is more cost-effective than if several NILM problems are implemented in different measurement points of the facility (e.g., one per main circuit) creating a hierarchical NILM system.

Regarding the features used in the load identification, low resolution features are highly desirable, because they can be computed from low-resolution electrical data gathered by cost-effective meters. In the last years, the NILM community have carried out significant progress in this area, focusing their efforts on approaches that rely on a single meter and use low-resolution data for the load identification, specially with the irruption of DL approaches.

Another aspect of the set-up cost is the computational complexity of the NILM solutions. High complexity algorithms entail the installation of special hardware, if the disaggregation is made *in situ*, or a connection with *ad hoc* web service in charge of the disaggregation [71], where all the required calculations are delegated to a dedicated server. In terms of cost-effectiveness, the second option is the cheapest one and, in addition, the use of a web service allows the user to consult their disaggregated energy consumption from anywhere and any device connected to Internet.

- *Unsupervised approaches.* Apart from the economic aspect, NILM systems involve costs in time of implementation, which are associated with the training time

needed to adapt the NILM approach to the facility to be analyzed. At this regard, purely unsupervised models are preferred, since they do not require any training data from the appliance-specific consumptions. The aforementioned unsupervised methods are far from supervised methods in terms of accuracy and most of them have a high computational complexity. For these reasons unsupervised methods with low complexity and trained with low resolution data are an open issue within NILM research.

- *Transferability.* Most of the NILM methods mentioned before assume that appliance signatures are always the same as the ones used during their training process, considering that the appliance-specific loads are independent of the context factors. However, context factors could bias the resulting models and compromise the performance of NILM systems when they are faced to “unseen” facilities, or when an important context factor varies significantly. In [111], authors divided these context factors in two categories:
 - *External environmental factors.* Those variables that do not directly affect the appliance’s operation, but modulate the usage of the appliances by the users. Weather conditions, seasonality-related attributes or changes in the normal occupancy of the facility are examples of these factors.
 - *Appliance-specific factors.* Factors that directly affect the signatures of the appliances. These factors are related to the operational modes of the appliances and the differences in the internal electrical and electronics elements used in the different designs of a same kind of device. Hence, modern devices integrate new operational modes in order to save energy (“eco” and “sleep” modes) that lead to new dynamics in their signatures, making them behave as *different* appliances.

One of the premises of DNN-based NILM models is that with enough training data from multiple possible contexts, DNN models could reach accurate load identification in unseen scenarios [71], but in practice the extremely complex dynamics induced by changes in the context variables and the lack of proper training datasets make that idea unfeasible. Although, in [109] it was pointed out that DNN-based models with less complexity may improve the overall transferability of DNN-based NILM models and in [91, 111] authors have suggested a method capable of retraining the models when a change in the context is detected, the transferability should be further studied in order to enhance the massive deployment of NILM systems.

- *NILM on non-residential facilities.* The idea of NILM could be applied to any electrical installation in which there are more than one energy-consuming devices downstream of the measurement point, and a fine-grained analysis is required. However, most of the aforementioned methods are focused on residential facilities and only a few of them have studied the NILM on non-residential installations, such as large commercial buildings [46, 114, 115]. The energy efficiency of large buildings has attracted a lot of interest in the last decade, since approximately about the 11% of the total energy consumption in developed countries is spent on these facilities, and therefore techniques such as the NILM systems that improve their energy efficiency by means of detailed energy monitoring tools are highly desirable. Transferring NILM methods designed for households to large non-residential buildings is not straightforward, since the nature and dynamics of the devices in large buildings, the data acquisition system and the amount of plugged-in devices are often different from residential facilities. Additionally, there are very few free available datasets [115].

VISUAL ANALYTICS FOR NILM IN LARGE BUILDINGS

Although energy load monitoring in large buildings seeks the same target as the residential load monitoring, non-residential buildings present substantial differences with respect to households in most of the NILM framework's stages mentioned in section 2.1. In this chapter, a deeper explanation about these differences will be presented by comparing both contexts through an insightful exploratory analysis of two representative datasets from both scenarios. Finally, an interactive and unsupervised approach for NILM in large buildings, designed following the *visual analytics* (VA) methodology, will be suggested and evaluated in sections 3.2, 3.3 and 3.4.

3.1 Differences between residential and non-residential facilities

In this first section, the differences between residential and non-residential facilities, in terms of electric energy consumption, will be briefly studied. For a fair comparison, two different datasets from both scenarios will be explored, whose facilities and data are representative of each scenario, being as close as possible to the typical data measured from electrical facilities in developed countries.

3.1.1 Datasets

In the NILM literature, there is a large number of available datasets [1, 112, 116, 117] that include consumptions of single-family residences. Most of them were collected from houses in developed countries, recording abundant data from main and individual consumptions in order to facilitate the training of supervised models, and their sampling periods usually range from 1 second to 1 minute (low-resolution).

On the other hand, the spectrum of available non-residential datasets is considerably smaller, and the few public datasets [114, 115] record a large number of individual consumptions with a small number of samples each, or their samples have been synthetically generated, as in the case of [114]. This severely limits the application of modern approaches (e.g., DNN-based models) which require a large amount of training data.

In the case of residential installations, the *UK-DALE repository* [1] could be considered as a representative dataset because:

- It gathers a large enough amount of samples from five different houses, with more than four years of samples in some of the recorded houses. Additionally, the energy consumption of the most common appliances in households (e.g., fridge, washing machine or electric oven) were exhaustively recorded along all houses.
- Data acquisition systems used for measuring both the whole-house energy demand and the demand of each individual appliance are recorded by low cost, single channel, low-frequency meters, similar to the meters commonly installed in real residential installations.
- The low-frequency data approach is aligned with the trend of affordable NILM systems, as it was explained in 2.5. Moreover, a low-resolution dataset reduces the data storage requirements for our experiments and the final deployment.

In table 3.1, a detailed view of the main features of the UK-DALE dataset is shown. Note that each house has its own context, including different energy-related features: number of occupants, year of construction and average energy consumption per day. Regarding the energy measurement, each house has different levels of detail in their measurements, being the *House 1* the most representative, since about 80% of the house's

3.1. DIFFERENCES BETWEEN RESIDENTIAL AND NON-RESIDENTIAL FACILITIES

total energy is monitored by measuring 52 individual loads (54 meters - 2 main meters) during more than four years.

In [1], the authors described the deployed data acquisition system for *House 1*, where the whole-house energy demand is measured using a PC sound card as a A/D converter, and setting the sample period to 1 s. On the other hand, appliance-specific consumptions are recorded by commercial plug-in individual appliance monitors (IAMs), which provide samples with a sample period of 6 s.

Table 3.1: General features of UK-DALE dataset extracted from [1].

House	1	2	3	4	5
Year of construction	1905	1900	-	1935	2009
Energy improvements	solar thermal, loft insulation, solid wall insulation, double glazing	cavity wall insulation, double glazing	-	loft insulation, double glazing	-
Number of occupants	4	2		2	3
Total number of meters	54	20	5	6	26
Number of main (whole-house) meters	2	2	1	1	2
Sample rate of main meters	16 kHz, 1 Hz, $\frac{1}{6}$ Hz (6s)	16 kHz, $\frac{1}{6}$ Hz (6s)	1 Hz, $\frac{1}{6}$ Hz (6s)	$\frac{1}{6}$ Hz (6s)	16 kHz, 1 Hz, $\frac{1}{6}$ Hz (6s)
Date of first sample	2012-11-09	2013-02-17	2013-02-27	2013-03-09	2014-06-29
Date of last sample	2017-04-26	2013-10-10	2013-04-08	2013-10-01	2014-11-13
Average main apparent energy kWh per day	8.90	8.00	12.35	10.24	17.56
Proportion of energy sub-metered	0.80	0.68	0.19	0.28	0.79
Date of last sample	2017-04-26	2013-10-10	2013-04-08	2013-10-01	2014-11-13

On the other hand, the *Hospital de León dataset* managed by the *SUPPRESS* research group [118] is a representative example for non-residential load analysis. The submetering architectures installed in large buildings, such as hospitals, are designed to record the main systems of the hospital with the aim of monitoring the energy consumption of the facility and preventing failures or malfunctions.

The hospital of the city of León is an example of this kind of facilities, where the energy consumptions along all the buildings and systems that constitute the hospital facility are monitored by means of powerful meters, which accurately record energy-related variables, such as currents, voltages, powers and power quality measurements (e.g., electrical harmonics and total harmonic distortion). The implemented submetering

architecture in the hospital is shown in Figure 3.1, where the measuring points are hierarchically distributed in three levels:

1. *Acometida* meter. It is connected at the supply connection of the Hospital, and records the whole-facility energy demand.
2. Meters installed in each of the two *general low-voltage panels* (Cuadro General de Baja Tensión, CGBT as its Spanish acronym). These two panels are directly connected to the supply connection, and they are in charge of distributing the energy along the different main circuits within the hospital. The meters installed in the CGBTs monitor all downstream loads on each of the main circuits.
3. Meters installed in each of the main circuits. These meters are in the lowest level of the submetering architecture and monitor important systems of the hospital, such as the elevators, the kitchen or the thermal comfort systems.

In the diagram shown in the figure, each measuring point is labeled with its meter ID, which corresponds to the specific model of meter from the *Schneider Electric* power monitoring catalog of meters [119]. Table 3.2 indicates in detail the model and references of the meters used for measuring each node of the submetering network.

Meter ID	Name	Reference
ION7650	PowerLogic ION7550	[120]
PM5110	PowerLogic PM5000	[121]
ION7650	PowerLogic PM8000 series	[122]
PM5560	PM5560 Meter	[123]

Table 3.2: Meters used in the hospital’s submetering network.

At first glance, the biggest difference between both datasets is the data acquisition system, since the UK-DALE repository uses cost-effective, single function and low-frequency meters, while the data from the hospital are recorded by expensive, multifunction and high-frequency meters.

3.1. DIFFERENCES BETWEEN RESIDENTIAL AND NON-RESIDENTIAL FACILITIES

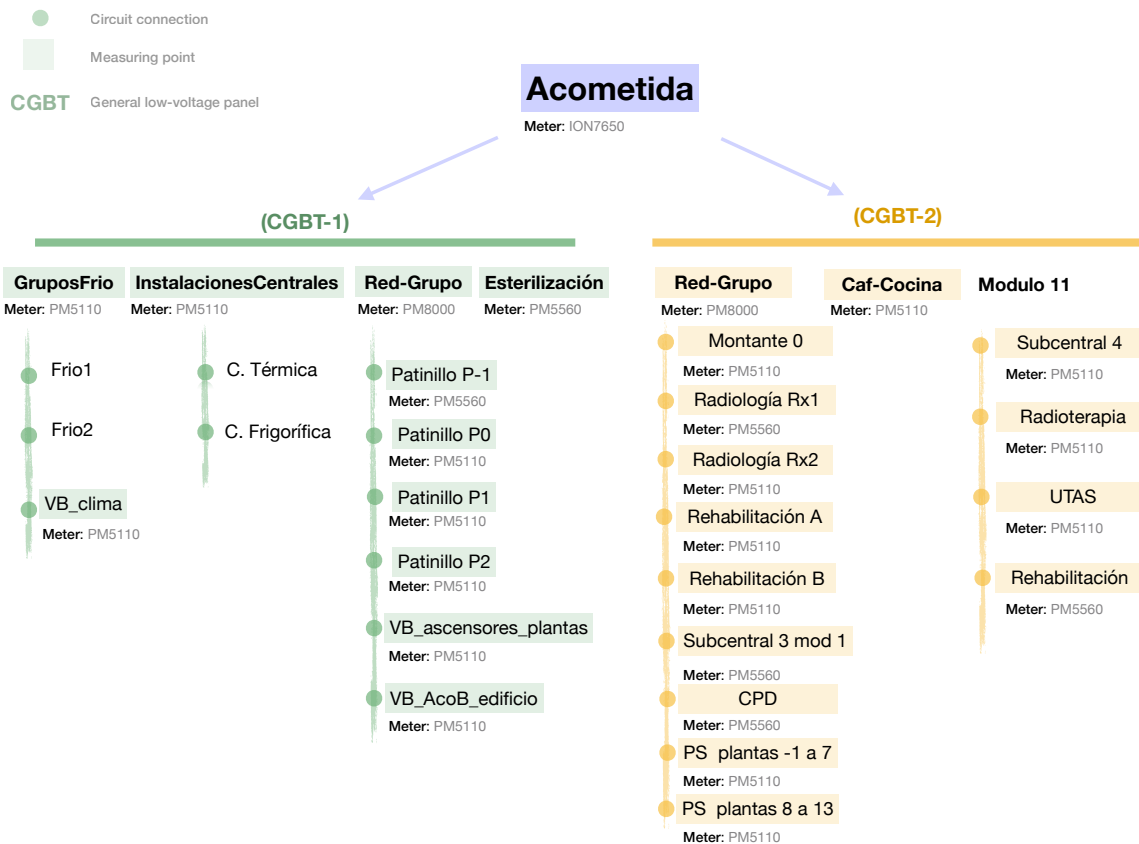


Figure 3.1: Submetering architecture of the Hospital de León dataset. The main circuits are represented as a vertical line, and the important connections are depicted as a point. All the measured connections are highlighted and identified by its label and the type of meter installed.

Apart from the type of meters, another difference is *where* the meters are placed. In the case of the hospital, there is a hierarchical set of meters permanently recording the consumption of the entire facility and the principal circuits, while in the case of the residential facilities, only the energy demand at the house supply connection (main consumption of the house) is measured permanently, and the individual consumptions are only recorded for training purposes.

These differences in the data acquisition systems condition the subsequent stages of the NILM framework, such as the features that can be extracted, or the energy disaggregation algorithms to be applied.

3.1.2 Nature of the consumptions in residential and non-residential facilities

In order to present a further comparison between the residential and non-residential energy demand data, an exploratory analysis of the energy consumption nature in the aforementioned datasets will be presented.

Table 3.3: General details of the CGBT2.Red-Grupo subset from the *Hospital de León dataset*.

Meter	Sample rate	No. samples	Description
Acometida.Acometida	1 min	519278	Meter connected to the supply connection
CGBT-2.Red-Grupo	1 min	519278	Consumptions from entire wing of the hospital
CGBT-2.Montante0	1 min	507353	All consumptions from lifts
Radiologia1	1 min	507305	Consumptions from X-ray room 1
Radiologia2	1 min	507354	Consumptions from X-ray room 2
RehabilitacionA	1 min	507355	Consumptions from rehab facilities A
RehabilitacionB	1 min	507307	Consumptions from rehab facilities B
Subcentral3	1 min	507353	Unknown
CPD	1 min	507353	Consumptions of the server and data center
Plantas_2-7	1 min	507353	Consumptions of the hospitalization floor from 2 to 7
Plantas_8-13	1 min	507355	Consumptions of the hospitalization floor from 8 to 13

3.1. DIFFERENCES BETWEEN RESIDENTIAL AND NON-RESIDENTIAL FACILITIES

In the case of the hospital data, only the consumption in the circuit Red-Grupo from the CGBT-2 and its downstream meters will be considered as our main consumption and individual consumptions, respectively. In Table 3.3, the individual nodes from CGBT-2.Red-Grupo are summarized, indicating the sample period, the number of samples and a brief description of each node. All data from CGBT-2.Red-Grupo have been subsampled to a sample period of one minute in order to reduce the amount of data stored in hospital databases.

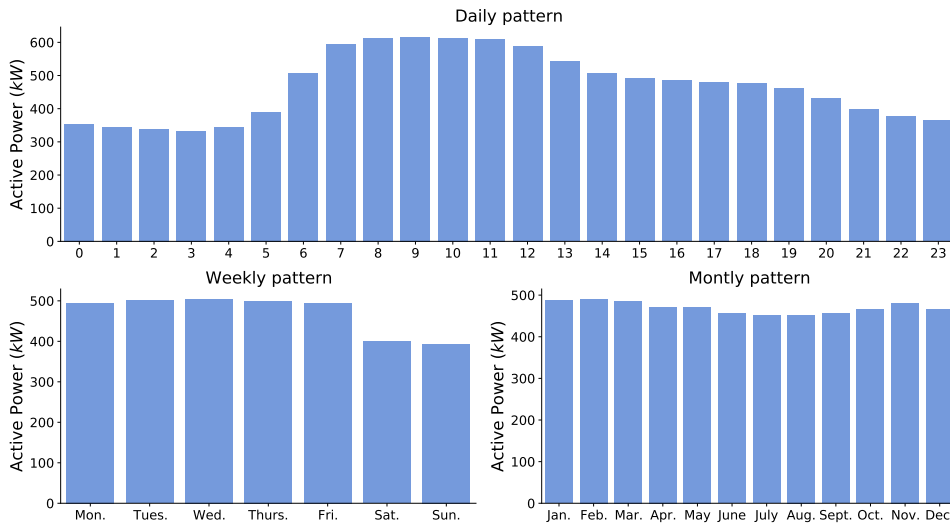


Figure 3.2: Average energy demand of the node CGBT-2.Red-Grupo grouped by different time scales (months, days of week and hours).

If we focus the analysis in the average of the main consumptions at different time scales, as in Figure 3.2, the daily pattern of CGBT-2.Red-Grupo from the hospital shows a daily hump shape, which is common in public service buildings whose activity is concentrated in the central hours of the day. In this kind of buildings, it is also common to observe a decline in the occupancy and activity during the weekend, which is reflected in the weekly view of the figure. Regarding the seasonality, the energy consumption in CGBT-2.Red-Grupo remains constant throughout the year except for a slight decrease in consumption during summer, due to a reduction in the occupancy and in the usage of the thermal comfort systems.

The same analysis for the House 1 of the UK-DALE repository is presented in Figure 3.3. At first glance, one can observe substantial differences in the daily profile, where the energy demand presents two humps instead a single hump. Daily energy demand in

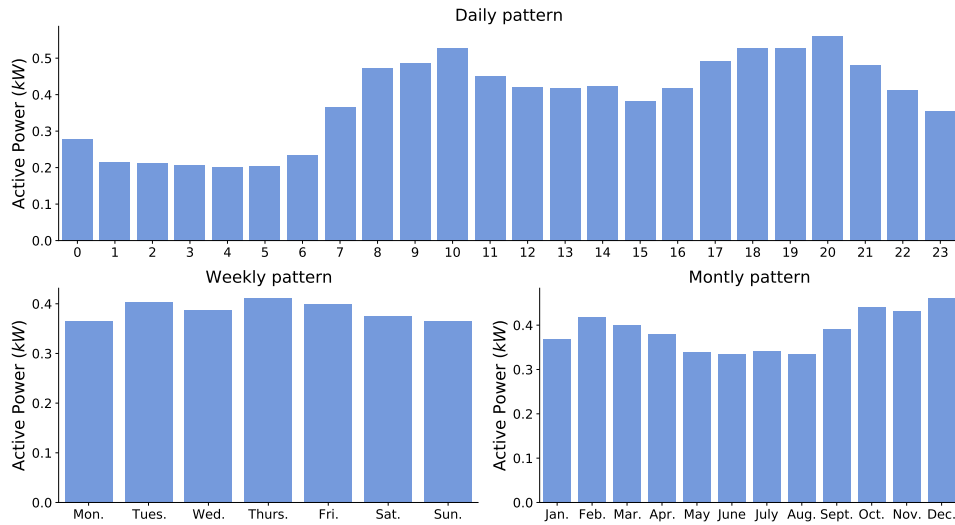


Figure 3.3: Average energy demand of House 1 from UK-DALE repository grouped by different time scales (months, days of week and hours).

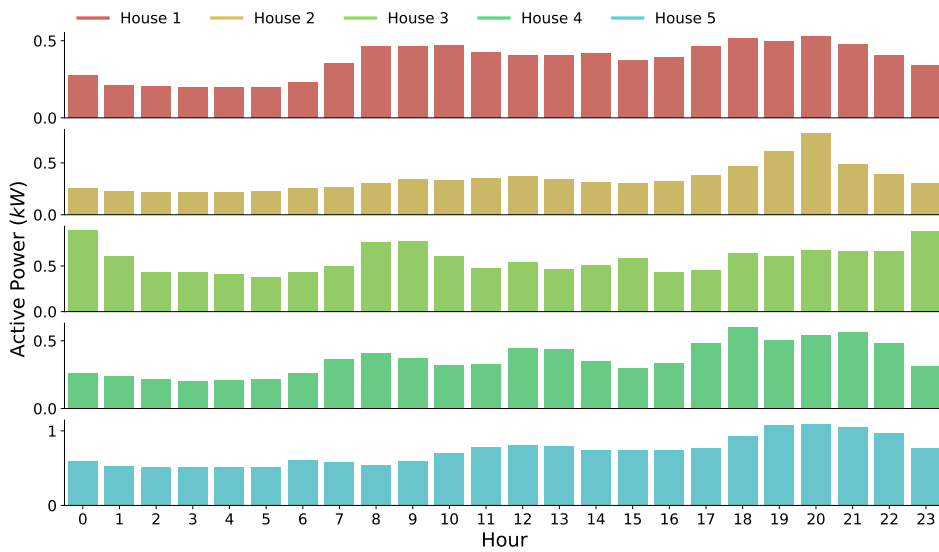


Figure 3.4: Daily profile of all the houses from UK-DALE repository.

houses strongly depends on the behavior and the consumption patterns of the occupants of the house, resulting in different daily profiles per house, as it is shown in Figure 3.4.

The house's weekly profile of Figure 3.3 is more stable than in the case of the hospital, being the differences in consumption between the weekend and workdays less pronounced. Additionally, the house's yearly profile shows a strong seasonality due to the energy demand of heating systems.

3.1. DIFFERENCES BETWEEN RESIDENTIAL AND NON-RESIDENTIAL FACILITIES

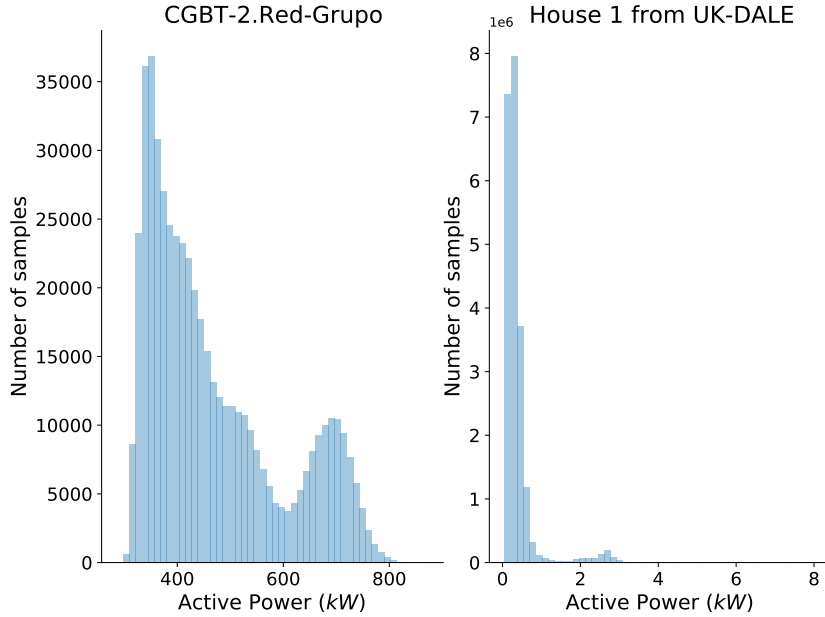


Figure 3.5: Histograms of main consumption of CGBT-2.Red-Grupo and House 1 from UK-DALE.

Looking at the histograms of the CGBT-2.Red-Grupo and House 1 consumptions, shown in Figure 3.5, we can observe that there are no recorded samples without energy demand in the hospital, while, in the case of House 1, the largest probability density is placed around zero. This indicates that the energy demand of houses is mostly spent in a few appliances of type I or II, while the systems connected to the electrical network of the hospital remain active during long periods of time (e.g. lights or thermal comfort systems), which can be considered as appliances of type IV.

This intuition is further analyzed in Figures 3.6 and 3.7, where an overview about the energy demand from the individual consumptions of the two facilities can be seen.

In the daily energy demand profiles shown in Figure 3.6a, the CGBT-2.Red-Grupo individual nodes have a *continuous energy demand*, since all of them aggregate consumptions from more than one device. Despite of this agglutination of devices, we consider these nodes as the individual consumptions to be disaggregated, since recording separately each device in a complex facility as the hospital would be unfeasible and extremely expensive.

Most of the individual consumptions in Figure 3.6a show a strong daily seasonality, except node CPD, which records the consumption of the servers' room and its activity

during the day remains constant. Daily seasonality causes differences between daylight and night consumptions, which are also highlighted in their histograms of Figure 3.6b. In addition to the daily seasonality, the profiles of some individual nodes present stationary and residual consumptions due to devices that remain active during long periods of times, such as the short-term and periodic pattern during the night in Radiologia1. This consumption is translated into pronounced spikes in its histogram of Figure 3.6b.

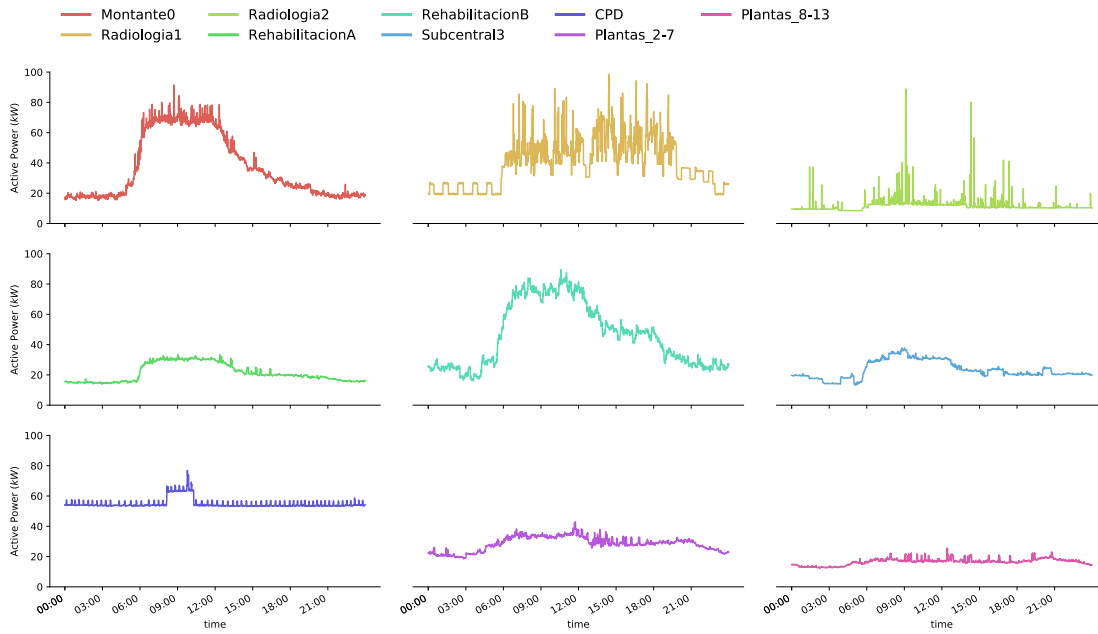
In the case of the individual consumptions of House 1 shown in Figure 3.7, only the five appliances with the largest impact in the main consumption (individual appliances with the largest number of activations) [42] are presented. The appliance daily profiles are more sparse than the individual nodes from hospital, since they are appliances of type I (e.g., fridge, kettle or microwave) and type II (e.g., washing machine or dish washer), as it is shown in Figure 3.7a. The sparse states of the appliances are captured in ranges of high density in their histograms of Figure 3.7b.

Bearing this in mind, we can conclude that the nature of the consumptions from the two facilities notably differs, being the data from the hospital individual nodes more structured, continuously active, stationary and dependent on the occupancy and activity of the building, while the energy demand data of the residential facility are sparse and dependent on the behavior of the people living in the house and their consumption patterns.

The nature of the hospital consumptions makes the event-based feature extraction difficult, since identifying events in a smooth and continuous waveform is not straightforward. This limits the approaches in the state of the art (see chapter 2) that rely on event detection stages, hindering their feature extraction, and therefore the subsequent energy disaggregation.

3.1. DIFFERENCES BETWEEN RESIDENTIAL AND NON-RESIDENTIAL FACILITIES

a) Individual consumptions



b) Histograms of individual consumptions

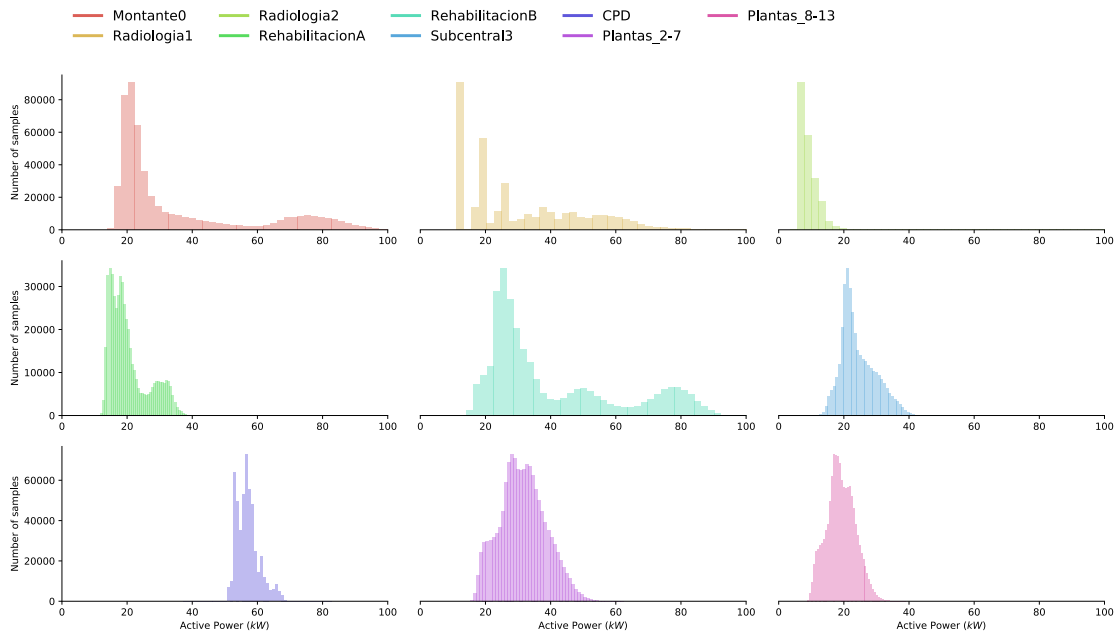



Figure 3.6: Intuition about the individual consumptions of the node CGBT-2. Red-Grupo. a) Energy demand of the individual nodes and b) histograms of the individual nodes.

a) Individual consumptions



b) Histograms of individual consumptions

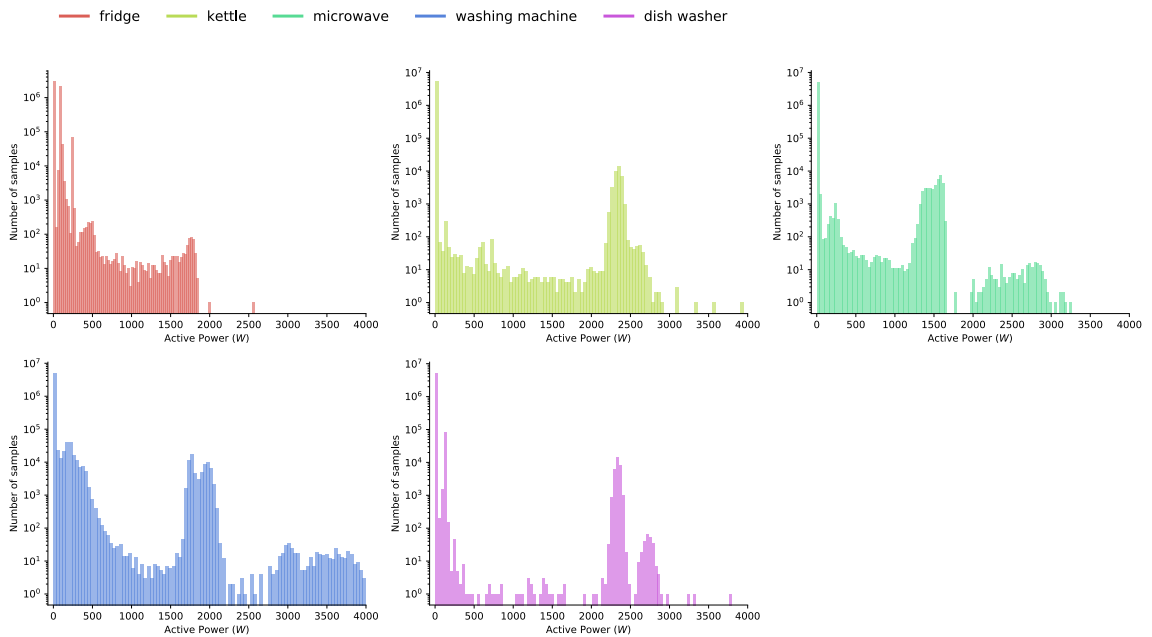


Figure 3.7: Intuition about the individual consumptions of the House 1 from UK-DALE repository. a) Energy demand of the individual appliances and b) histograms of the appliance-specific consumptions. Note that these histograms are in logarithmic scale and most of the recorded samples are zero.

3.1.3 Visual analytics as a first approach to NILM in large building

Considering that residential and non-residential facilities are different contexts and that the nature of their consumptions is also different, the performance of the state-of-art methods designed for residential facilities could be affected when they are applied to large buildings, especially those which rely on event-based features. In this chapter, we will introduce a novel approach for NILM in large non-residential buildings based on the *visual analytics* (VA) paradigm with a high exploratory component. A VA perspective could be a convenient first approximation in a scenario as NILM in large buildings, from which analysts and maintenance staff of the building can insightfully increase their general knowledge of the facility and familiarize themselves with patterns of downstream consumptions, before addressing more complex analysis (e.g., HMM or deep auto-encoders).

In this first approach, a key factor is the previous knowledge of the maintenance staff which, according to our problem formulation mentioned in section 1.2, is modeled as K . This expert knowledge is another difference from the NILM on residential facilities, since in large and complex electrical installations, such as the hospital, there are often experimented staff with knowledge accumulated over years of plant operation that in residential facilities do not exist. Hence, in the first stages of the analysis, this expert knowledge can be harnessed to enhance results and conclusions.

In the VA diagram shown in Figure 3.8, the expert knowledge K is exploited to prompt new hypotheses that modulate both the NILM operation D and the data visualization operation V by means of their parameters S_D and S_V , respectively.

In subsequent sections of this chapter, we shall focus our attention on adapting NILM techniques to be interactively visualized by means of data visualization techniques V so that an insightful view I of the resulting disaggregation $\mathbf{p}(t)$ is provided to the user. In addition, we will explore powerful mechanisms of interaction which ensure that the operations V and D can be modulated by users by modifying their parameters S_V and S_D according to new user's hypotheses.

A first stage of our approach is to define a suitable energy disaggregation operation D in our VA formulation. Taking into account that our VA approach is suggested for an exploratory and preliminary analysis, the energy disaggregation technique must not

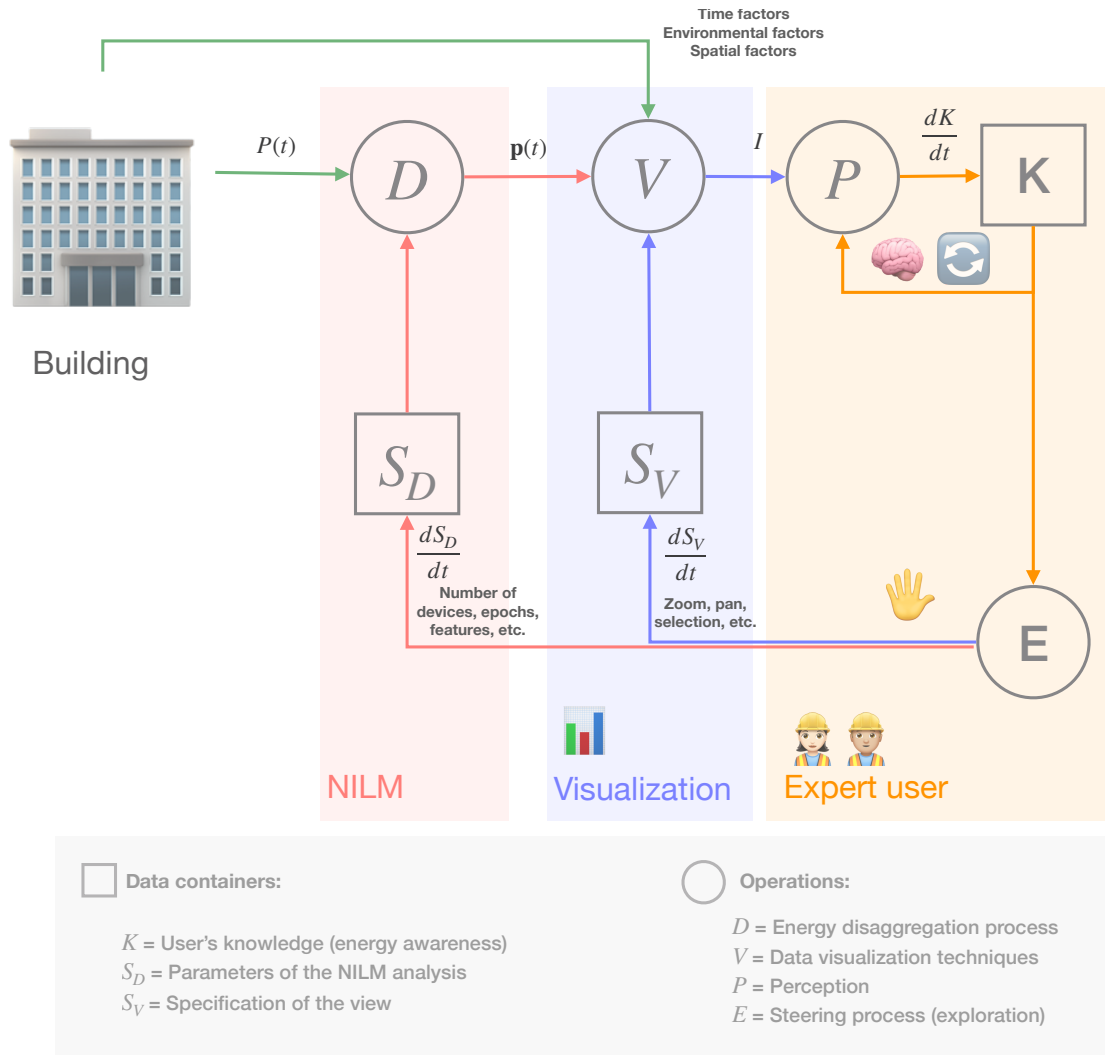


Figure 3.8: Diagram of our visual analytics perspective of NILM particularized for large buildings.

depend on information about electrical installation (i.e., number of circuits, features of downstream systems or data acquisition equipment) or the availability of data from the individual nodes in the training stage. For that reason, unsupervised models that only need the main consumption in order to decompose the main energy demand of the facility are convenient for easily reproducible approaches. Selected energy disaggregation techniques should also have a *low complexity*, being as simple as possible in order to reduce the latencies in producing a new result after a change in their parameters S_D .

Some of the unsupervised NILM models in the literature, such as HMM-based models, take a long time to process a new disaggregation, and therefore do not satisfy our low-

complexity requirement. Taking this into account, among the energy disaggregation techniques described in section 2, only *blind source separation* (BSS) and *graph-based signal processing* (GSP) NILM approaches are suitable for our disaggregation technique D due to their low latencies.

Regarding the data visualization operation V of the diagram in Figure 3.8, it is important to note that data recorded from large buildings involve other *context attributes* (some of them known and measured and others unknown) apart from the electrical measurements, such as time factors (e.g., hour, day, week or year), spatial factors (e.g., buildings, floors or subsystems of the building), environmental factors (e.g., temperature, humidity or occupancy) or other relevant attributes associated with systems that are considered critical for the normal operation of the building. The fact that the main consumption of the building is considerably affected by all these factors, suggests that the energy monitoring in large buildings could be considered as a *multiway* problem.

In our VA approach all these factors should be taken into account, and the operation V should take advantage of them to produce intuitive views, where an expert user can discover hidden insightful correlations between all these factors and the estimated individual consumptions from the energy disaggregation. In this scenario, several works [124, 125] have suggested that these correlations can be found through well designed visual encodings capable of transforming information into appropriate representations. Most of these representations are based on static 2D maps which are only able to display a few factors, but our VA, in the same way as in the operation D , should provide users with mechanisms that allow to modify the view according to user's hypotheses, resulting in a "live" sequence of views instead of a static view. In addition, the transition from one view to a new view should be fluid enough, so that users feel an immediate feedback increasing the confidence in the obtained correlations [126].

3.2 Blind Source Separation approaches for NILM in large buildings

Looking back to the initial definition of energy disaggregation of our problem formulation in section 1.2:

$$P(t) = p_1(t) + p_2(t) + \dots + p_m(t)$$

one can consider this definition as a single-channel separation problem, where the total consumption $P(t)$ is the result of a mixture of different sources, in this case the individual consumptions $p_m(t)$. Hence, the nature of the NILM problem is closely related to the *cocktail-party* problem, where the target is to identify or to separate what is the individual speech of each speaker in a party with a mixture of voices from several interlocutors [127].

In [101], authors addressed this unsupervised perspective of NILM as a single-channel problem, using *sparse coding* algorithms to separate the main energy demand in residential facilities into individual sources. Moreover, they introduced constrains in the computation of the sources that adapt the sparse coding algorithm to the nature of the NILM task, such as positivity or sparsity constraints on the obtained sources.

Following this paradigm of NILM as a single-channel separation, we suggest an energy disaggregation approach based on the *Non-negative Matrix Factorization* (NMF) [5] algorithm for our VA approach. In the following subsections, the application of NMF to the energy disaggregation problem and why it is convenient for our VA formulation will be explained in detail.

3.2.1 Non-negative Matrix Factorization formulation

In [5, 128, 129], authors defined NMF as a decomposition of a non-negative input matrix of observations. In the problem of single-channel decomposition, the observations matrix is made of different samples from the same mixture of sources. This matrix of positive observations is defined as \mathbf{V} and the resulting decomposition is the result of the matrix multiplication of a *matrix of bases* \mathbf{W} and a *matrix of activations* \mathbf{H} :

$$\mathbf{V} \approx \mathbf{W}\mathbf{H} \tag{3.1}$$

Matrix \mathbf{V} could be defined by M -dimensional column vectors \mathbf{v}_j whose elements are non-negative so that:

$$\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N] \in \mathbb{R}_+^{M \times N} \tag{3.2}$$

where vectors $\mathbf{v}_j, j = 1 \dots N$ form a set of N observations of the mixture to be separated.

In the same way, the matrix \mathbf{W} could be defined as

$$\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_L] \in \mathbb{R}_+^{M \times L} \quad (3.3)$$

being \mathbf{w}_i its i -th column vector, and the matrix \mathbf{H} can also be expressed as

$$\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_L]^T \in \mathbb{R}_+^{L \times N} \quad (3.4)$$

where \mathbf{h}_α^T is the α -th row of the matrix \mathbf{H} . Taking matrices \mathbf{W} and \mathbf{H} by columns and rows, respectively, NMF decomposition can be reformulated as:

$$\mathbf{v}_j = \sum_{\alpha=1}^L \mathbf{w}_\alpha h_{\alpha j} \quad (3.5)$$

Hence, the observation vector \mathbf{v}_j is the result of a linear combination of the columns in \mathbf{W} , whose coefficients are the elements of the columns in \mathbf{H} . In other words, all input observations are defined as a weighted sum of non-negative basis vectors.

The positive nature of the basis vectors \mathbf{w}_i suggests that they can be considered as “*the parts of a whole*”, which can be combined in order to reconstruct an example of the observed mixture. The importance of the bases \mathbf{w}_i in this combination is determined by the elements in matrix \mathbf{H} , being the activation of the α -th basis along the observations, the elements of the α -th row in \mathbf{H} .

The resulting decomposition is an insightful, positive, parts-based representation of the observations. In [5], this parts-based representation was exemplified applying NMF to images of human faces so that the relevant parts of a human face (noses, mouths, eyebrows) were identified in the obtained bases of \mathbf{W} , as it is shown in Figure 3.9.

3.2.2 Complexity of NMF

In [130], authors showed that an exact NMF decomposition, such as $\mathbf{V} = \mathbf{WH}$ is a NP-hard problem. Taking our low-complexity requirement into account, we dismiss exact NMF approaches, and the approximated decomposition in equation (3.1) is alternately applied. Another drawback of this basic NMF approach is having to set the number of bases L




Figure 3.9: NMF decomposition of faces from [5]. Matrix on the left represents the matrix of bases \mathbf{W} and the matrix on the right represents the matrix of activations \mathbf{H}

in \mathbf{W} , since the accuracy of the reconstruction \mathbf{WH} is strongly related to this parameter. L can be effectively selected according to the performance of the NMF decomposition, measuring how close is the reconstruction \mathbf{WH} to the observation matrix \mathbf{V} by means of a reconstruction error metric, such as the *mean squared error* (MSE), and then selecting the number of components with the best reconstruction error.

In complex scenarios, as it could be the application of NMF to data from large facilities, this method could lead to a large number of components, making the interpretation of the results more complicated. In our approach, following the VA problem formulation, we propose a trade-off solution, where the user is able to set L according to his previous experience, his analysis requirements and the feedback obtained in previous results (new hypothesis generation of VA approach). Thus, L is one of the settings S_D by which the user can modulate the NILM operation D in our formulation of Figure 3.8.

3.2.3 Objective function and SVD-based initialization

Once the NMF decomposition is established, the matrices \mathbf{W} and \mathbf{H} are estimated by minimizing an objective function, which is typically set to a similarity measurement [5] between \mathbf{V} and its reconstruction \mathbf{WH} , such as:

$$\begin{aligned}
 J(\mathbf{W}, \mathbf{H}) &= \frac{1}{2} \sum_{i=1}^M \sum_{j=1}^N (V_{ij} - (WH)_{ij})^2 \\
 (\mathbf{W}^*, \mathbf{H}^*) &= \arg \min_{\mathbf{W}, \mathbf{H}} J(\mathbf{W}, \mathbf{H})
 \end{aligned} \tag{3.6}$$

where the optimization is subject to $W_{i\alpha} \geq 0$, $H_{\alpha j} \geq 0$, $\forall i, j, \alpha$, resulting in a bounded-constrained optimization. This optimization problem is not convex with respect to \mathbf{W} and \mathbf{H} at the same time, but it is convex with respect to \mathbf{W} or \mathbf{H} separately [131]. Hence, a coordinate gradient descent algorithm [132] could be applied, where the gradients are projected to the feasible positive region, satisfying the rule:

$$\begin{aligned}
 \mathbf{W}^{k+1} &= \max(0, \mathbf{W}^k - \beta_k \nabla J(\mathbf{W}^k, \mathbf{H}^k)) \\
 \mathbf{H}^{k+1} &= \max(0, \mathbf{H}^k - \beta_k \nabla J(\mathbf{W}^k, \mathbf{H}^k))
 \end{aligned} \tag{3.7}$$

where β_k is the step size for each iteration k . Resulting gradient descent could lead matrices \mathbf{W}^k and \mathbf{H}^k to local minima and the overall optimization strongly depends on the initialization of matrices \mathbf{W} and \mathbf{H} .

The first approaches of NMF initialized \mathbf{W} and \mathbf{H} with positive random values, which led to different solutions in each training process. Subsequent NMF approaches have proposed basic low rank decompositions of \mathbf{V} as convenient initializations for \mathbf{W} and \mathbf{H} . In [133], a *nonnegative double singular value decomposition (nndsud)* is suggested as an initialization method for matrices \mathbf{W} and \mathbf{H} . Authors showed that nndsud initialization speeds up the gradient-based optimization of NMF without any randomization. The absence of random process makes svd-based initialization suitable for VA purposes, since the resulting NMF decomposition after a retraining triggered by an user interaction will be always driven towards the same local minimum.

3.2.4 Sparseness of NMF decomposition

The aforementioned NMF formulation can be seen as a low rank decomposition which is totally additive, because the reconstruction of the observations (3.5) lacks cancellations by negative components. This positive parts-based representation is naturally interpretable, since cancellations between bases will be more confusing for users.

In addition to the absence of cancellations, solutions of NMF computed by the optimization problem defined in (3.6) tend to be sparse as well. Coordinate gradient descent in (3.7) updates matrices \mathbf{W}^k and \mathbf{H}^k so that in each step of the NMF optimization the resulting estimations are closer to the observations, but if any element of these matrices falls outside of the positive region, it is immediately projected to zero due to the positive constraints, as it is visually represented in Figure 3.10.

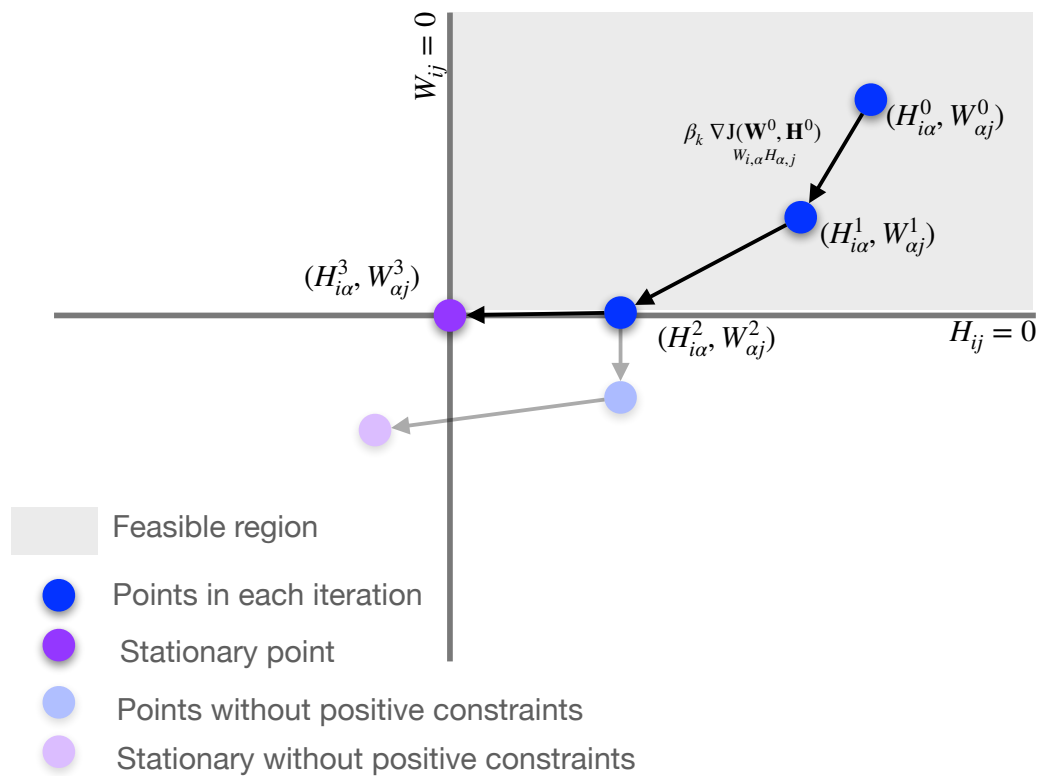


Figure 3.10: Intuition about the coordinate gradient descent of the NMF decomposition. Possible NMF solutions rely on the feasible positive regions.

This highlights that the elements of \mathbf{W} and \mathbf{H} in NMF solutions may be clipped to the positivity borders ($W_{i\alpha} = 0$, $H_{i\alpha} = 0$), making the resulting parts-based representation sparse. Therefore, NMF decomposes an input matrix \mathbf{V} into positive parts that are only relevant in a few elements. This quality of NMF decomposition is strongly connected to the human mechanism of *learning objects by its parts*, and therefore makes the results naturally insightful and suitable for the subsequent data visualization and interaction approaches.

Despite in first works in NMF decompositions [128, 129], the authors demonstrated that NMF decompositions of real data yield sparse bases and activation coefficients, in [131] the authors integrate new terms in the cost function in order to ensure the sparsity of the results, such as the L_1 norm. In our work, we consider that the basic NMF leads to sparse enough parts-representation, and therefore no sparsity constraints are included in the optimization.

3.2.5 Adaptation of NMF to time series

Once the end-to-end process of the NMF decomposition is described, a methodology for applying it to the main consumption $P(t)$ should be defined. Transferring the concepts of NMF to time series involves composing an input matrix \mathbf{V} . In the aforementioned example of face decompositions, it is easy to define \mathbf{V} because data are intrinsically made of discrete samples of the problem (pictures of human faces), but in the case of time series problems, data involve continuous sequences, and therefore a method for breaking these time sequences into chunks is needed. Hence, we form matrix \mathbf{V} by breaking the sequence $\{P(t)\}$ up into N contiguous windows of length M :

$$\underbrace{\{P(0), P(1), \dots, P(M-1)\}}_{\mathbf{v}_1}, \dots, \underbrace{\{P((N-1)M), P((N-1)M+1), \dots, P(NM-1)\}}_{\mathbf{v}_N}$$

The N chunks obtained from the windowing process are arranged in the matrix \mathbf{V} as follows:

$$\mathbf{V} = \begin{pmatrix} P(0) & P(M) & \dots & P((N-1)M) \\ P(1) & P(M+1) & \dots & P((N-1)M+1) \\ \vdots & \vdots & \ddots & \vdots \\ P(M-1) & P(2M-1) & \dots & P(NM-1) \end{pmatrix} \quad (3.8)$$

The resulting input matrix and the windowing operation is defined by the window size M , and it directly conditions the outcome of the NMF decomposition.

As it was shown in section 3.1.2, the main energy demand in large buildings strongly depends on the human activity, and it is structured by days, weeks, seasons or years. Taking this into account, it seems appropriate to set M so that the resulting windows would be daily, weekly or yearly instances of the main consumption.

In addition, by setting M to one of these recognizable periods, a more interpretable decomposition will be achieved, as it will be explained in subsequent sections. The granularity of the input windows will determine the granularity of the NMF decomposition as well, and therefore M should be chosen according to the desired analysis. Thus, with annual windows, the patterns observed in the bases will be coarser than the patterns obtained with daily windows.

Once the NMF decomposition is computed, the matrix \mathbf{V} can be decomposed as:

$$\mathbf{V} = \mathbf{w}_1 \mathbf{h}_1^T + \mathbf{w}_2 \mathbf{h}_2^T + \dots + \mathbf{w}_L \mathbf{h}_L^T \quad (3.9)$$

where the outer product $\mathbf{w}_i \mathbf{h}_i^T \in \mathbb{R}^{M \times N}$ has the same units and dimensions as the observations in \mathbf{V} (e.g., kW in the case of the active power).

Thus, undoing the windowing operation on the product $\mathbf{w}_i \mathbf{h}_i^T$, a sequence of the same length as the initial sequence $\{P(t)\}$ is obtained, and it can be considered the L -th individual consumption \mathbf{p}_L estimated by the NMF decomposition. The individual consumptions $\mathbf{w}_i \mathbf{h}_i^T$ ¹, the basis consumptions \mathbf{w}_i and their activations \mathbf{h}_i^T are the output of the operation D in our VA scheme, and they will be the raw input information for the data visualization technique V .

3.3 Visual techniques in large buildings NILM

Once the operation D of our VA approach is defined, the next step is to represent the information about the facility gained from the NMF decomposition. For that purpose, an intuitive *data visualization* (DV) approach for the decomposition that, by tweaking its settings S_V , provides interaction paths for users is needed. An application with these functionalities offers the user natural mechanisms for raising new hypotheses, finding connections between the estimated decompositions and with context variables, or detecting anomalies in the facility.

Designing a suitable visualization for a particular dataset is not a straightforward process and it often requires a solid knowledge of the nature of the data to be represented and the inner mechanisms of the human cognitive system. During last decades, many

¹In subsequent sections we indistinctly refer to $\mathbf{w}_i \mathbf{h}_i^T$ as individual consumption, NMF factor or individual component.

works [134–137] have laid the bases of the DV field, enumerating the vast benefits that DV approaches bring to the data science and defining basic design principles and techniques of interactive DV models for different types of data.

These principles of design must be taken into account in order to define a suitable operation V and its interaction mechanism S_V (see Figure 3.8). In this section, we will explore the DV principles of design from an energy monitoring perspective to come up with a convenient interactive data visualization for the suggested NMF decomposition.

3.3.1 Principles of visual coding

A visual representation of complex information is the process of efficiently displaying multidimensional data in a 2D view, such as a screen or a page. In his books [138, 139], Edward R. Tufte defined some general principles that graphical representations must follow to come up with efficient representations of complex quantitative data or, as he called it: “*graphical excellence*”.

Displaying as much information as possible in the available space is the main design principle. Graphics often present artifacts that are plotted as ornamentation, with no other purpose than to beautify, and therefore they distract the user’s attention from the data elements. All those visual elements unrelated to the information under the study were denominated by Tufte as *chartjunk*.

In order to prevent chartjunk, Tufte suggested maximizing the ratio of total amount of ink used for data elements, called *data-ink ratio*.

$$\text{data-ink ratio} = \frac{\text{total amount of ink used for representing data}}{\text{total amount of ink used in the graphic}}$$

Maximizing the data-ink ratio prompts designers to good practices, such as avoiding redundancies, sharing ink between data elements and increasing the density of data elements. However, *how can we display huge multidimensional modern datasets without enlarging excessively the overall plot area?*

Tufte suggested the *small multiples* principle, by which the plot area of a graphic is split into small and condensed views of different poses from the same data, increasing the data-ink ratio and providing a comparative visual analysis between poses (see the

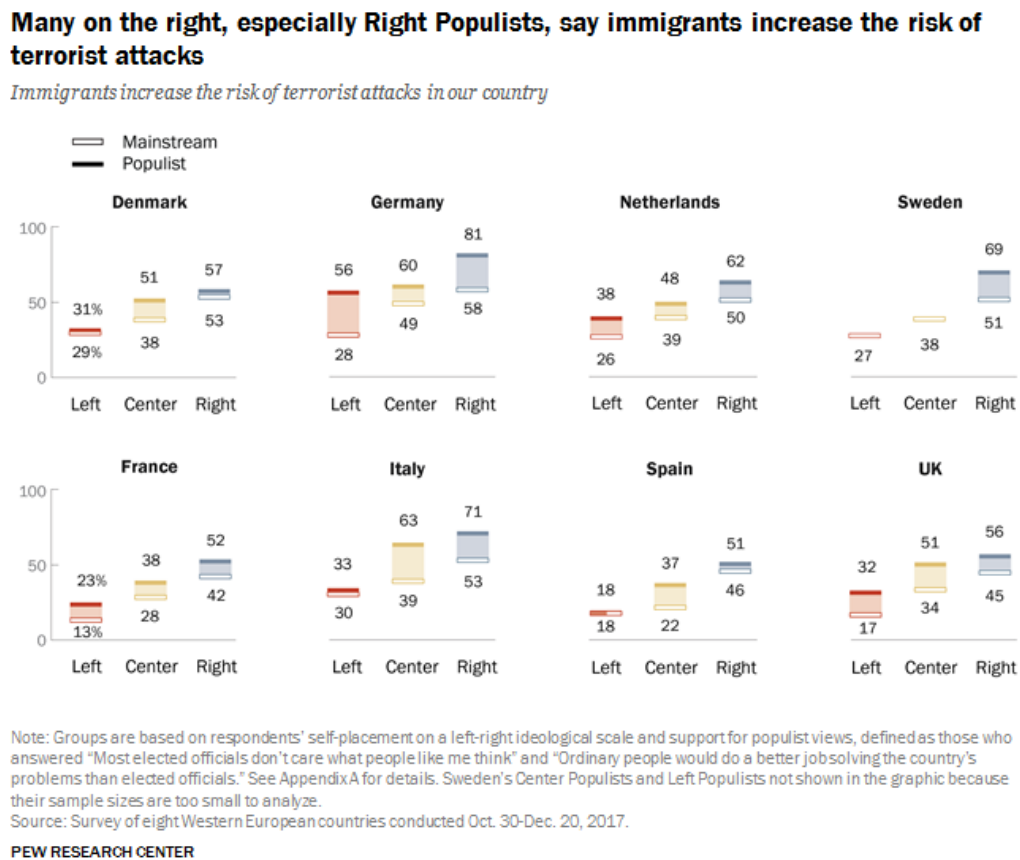


Figure 3.11: Example of *small multiples* from (<https://pewrsr.ch/3o4Suml>).

example of Figure 3.11). Small multiples approach can be easily translated to our NILM problem by, for example, plotting in a grid all the available individual consumptions together with the main consumption.

Other design principles related to the *types* of data variables to be represented have been studied. Hence, Tamara Munzner in her book [6] pointed out the importance of the type of data in designing good visualizations, and she divided the types of *attributes* into the categories shown in Figure 3.12, which are defined as follows:

- *Categorical*. Those data types which lack intrinsic order, such as types of appliances (e.g., kettle, fridge, washing machine, etc.).
- *Ordered*. Data types that have an intrinsic order. Examples of this category are attributes such as: temperature, energy demand or trouser's size. Ordered data can be also subdivided into:

- *Ordinal*. Those data attributes which are not completely continuous but show an explicit order. For instance, clothing size M is between sizes S and L .
- *Quantitative*. Numeric attributes that admit arithmetic operations such as the height of a person, active power, temperature etc.. Quantitative data types can be subdivided into two more categories: 1) *sequential* data, which show a homogeneous sequence from a minimum value to a maximum value; and 2) *diverging* data, where the data are centered in a value, from which two different sequences of values come out pointing to different directions.

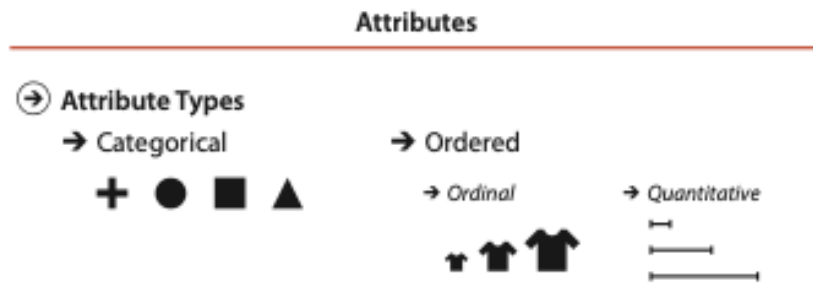


Figure 3.12: Data type taxonomy extracted from [6].




Figure 3.13: Combination of marks and channels.

Taking into account the Munzner's taxonomy, in our problem, only categorical (e.g., types of individual consumptions) and sequential attributes (e.g., energy-related vari-

ables and context factors) are involved, but *how can we visually encode these attributes efficiently in terms of human perception?*

In her book, Munzner suggests that appropriate visual encodings for attributes can be designed as a combination of two basic visual identities: *marks* and *channels*. Marks are basic visual elements that can be imputed to a record from data. For instance, a point is a zero-dimensional mark, a line is a 1D mark, whereas an area is a 2D mark. Position, size, color, orientation, hue, saturation and lightness are examples of channels, which visually define any mark independently of its dimensionality. In Figure 3.13, possible combinations of channels and marks are shown.

Designers have the responsibility to select an appropriate mark-channel combination for each attribute in order to guarantee the veracity of the visualization. In [6], Munzner discusses the *effectiveness* of channels and marks, dividing them into two categories: *what* and *how much* encodings. Inspired by this categorization, in Figure 3.14 a taxonomy of the channel-mark combinations ranked by effectiveness is shown.

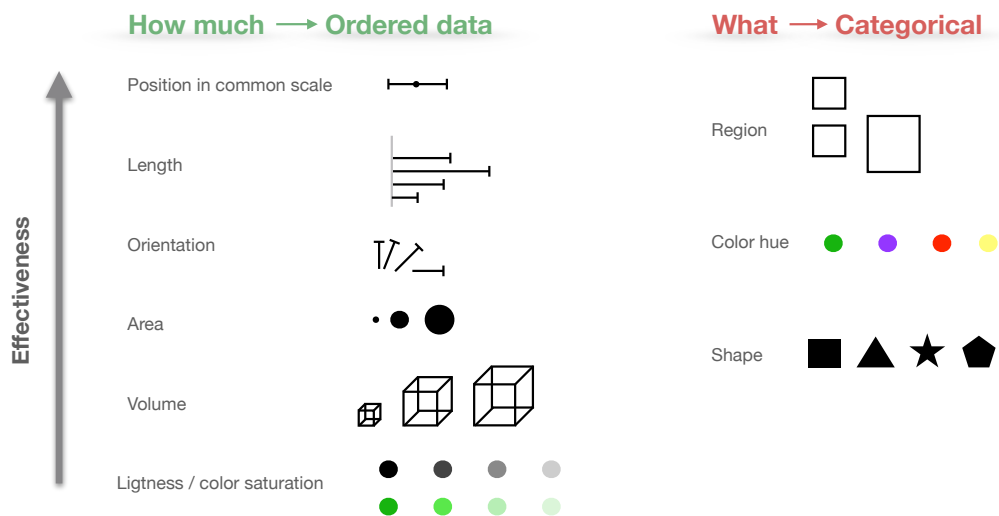


Figure 3.14: Taxonomy of channels effectiveness. Inspired by the discussion about the use of channels in [6].

In the specific case of our energy disaggregation problem, most attributes are sequential (e.g., temperatures, electrical variables or time variables), and the main purpose of our visualization is to compare individual consumptions among them or with context variables. *Therefore, our graphic design should include position and length for magnitude comparisons and for envisioning correlations.*

Additionally, if a categorization is needed (e.g., having different individual profiles in the same graphic), each of its items will be encoded by a different color hue or position. The principles of visual design will be revisited later in section 3.3.4, where the specific visual representations of our suggested VA approach are explained.

3.3.2 Principles of interaction

Human beings can intuitively interpret static views with high density of information, even in extremely condensed graphics. Our visual cortex allows us to easily modulate our analysis from a simple and comparative overview, to a fine-grained search in the view.

Although this *micro/macro reading capability*, as it was called by Tufte [139], enables the designer to display a large amount of data in the same view, it is only effective when data contain a limited number of samples and attributes. Thus, a condensed view of a city map is a successful example of a micro/macro reading, since only a limited number of attributes are involved (i.e., localization, street names, transport information) and their semantics is quite familiar to the users.

By contrast, in complex scenarios such as energy data from large buildings, where multiple energy-related and context attributes are involved, designing an effective micro/macro reading without losing information is not straightforward.

In these complex scenarios, the flexibility of the representation should be increased through *mechanisms of interaction* by which users can change the view by filtering the information, rearranging the elements or changing the channel-mark combinations (appearance). In this way, the user does not analyze just an static view I , but he is able to analyze a sequence of views $I(t)$ throughout the time of analysis t . The changes between contiguous views in time are triggered by tweaking the S_V parameters (sliders, zoom tools, etc.), according to our problem formulation of Figure 3.8.

Hence, a view $I(t)$ can be considered a point in the infinite space of the possible views defined by its settings S_V . In this space, the user can draw a trajectory through changes in S_V , that drives $I(t)$ according to his requirements. This enables to easily go from general views to more specific displays, obtaining a micro/macro reading in an interactive way.

The resulting human-in-loop analysis is tightly connected to the *sensemaking* process

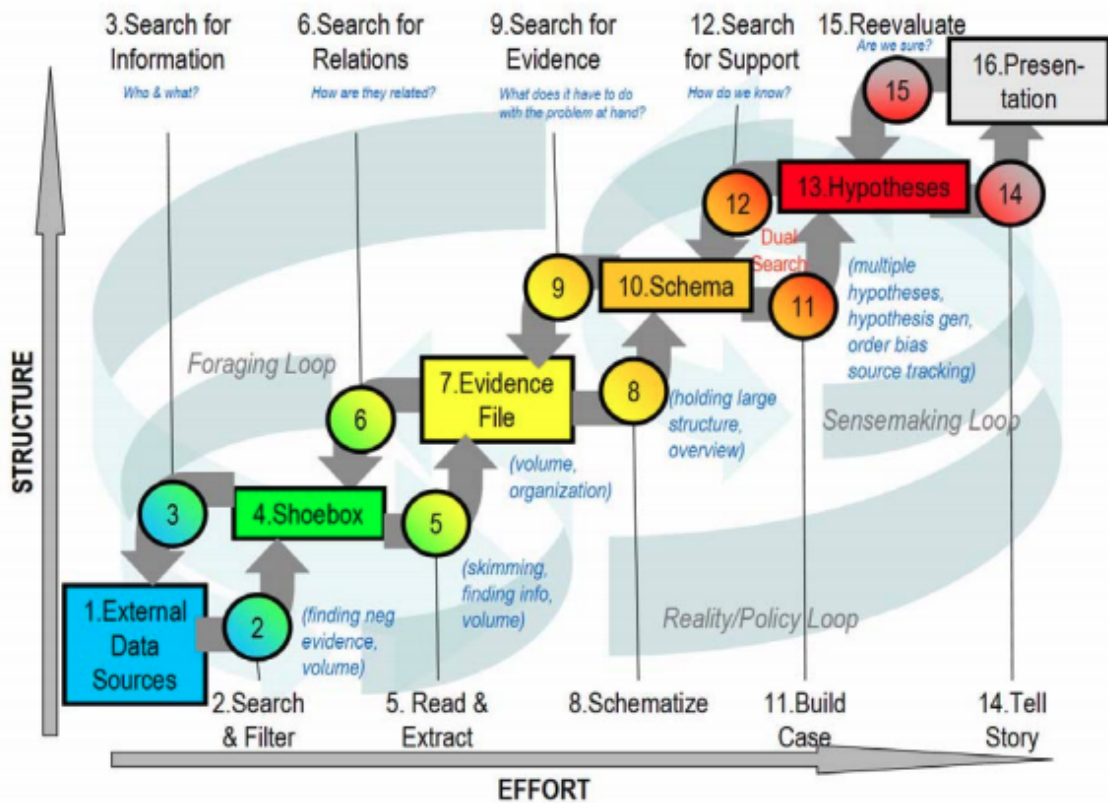


Figure 3.15: Sense making process extracted from [7].

of our cognitive system [7, 21] depicted in Figure 3.15, where forward and backward interactions in abstraction take place. Fluidity in these forward and backward passes is a limiting factor, so that fluid interaction mechanisms significantly increase the overall confidence in the approach and help to consolidate the gained knowledge [126].

At this regard, several works [140–142] have defined several best practices for the design of interaction pathways, classifying interaction mechanisms according to the effects caused in the data displayed. In [143], previous taxonomies of interaction techniques were reviewed and summarized in seven categories based on the user’s intentions:

1. *Select* mechanisms enable the user to highlight a set of visual items as a group of interest.
2. *Explore* interaction techniques provide the user with tools for changing the subset of data under analysis.

3. *Reconfigure* interaction mechanisms are those which change the spatial arrangement of the visual elements of the view.
4. *Encode* interaction mechanisms enable the user to change the visual encoding of data by modifications in the channel-mark combinations.
5. *Abstract/Elaborate* interaction techniques allow the user to increase or decrease the abstraction of the view. These mechanisms are aligned with the idea of micro/macro readings.
6. *Filter* methods configure the view so that all the data elements in the view must supply a set of conditions imposed by the user.
7. *Connect* interaction techniques enable users to underline elements, that have already plotted, and user considers that are associated between each other.

In our energy load monitoring application, as it was mentioned in section 3.1.3, we face a *multiway* problem where users may need to address the analysis by parts, gaining knowledge from subproblems. This division into subproblems can be easily addressed by means of *reconfiguration* or *filtering* interaction mechanisms. On-demand interaction techniques that change the level of abstraction are also suitable for energy load monitoring tools, where punctual failures or events must be examined in detail.

Additionally, if small multiples are integrated in our view to represent several time series of the same facility at once (different perspectives of the same system), *connecting* data through the views is also desirable (e.g., highlighting an event of interest in several plots or perspectives of different electrical variables).

3.3.3 The data cube paradigm

In order to enhance the energy load monitoring obtained from NILM methods in large buildings, we have suggested the VA approach shown in Figure 3.8, which immerses users in a human-machine loop interaction, where they must be able to formulate new hypotheses such as: when does this individual consumption take the highest value?; where is this individual consumption more relevant?; is there any correlation between this individual consumption and any of the context attributes?

Answering these questions implies providing with mechanisms of interaction between the user and both the graphic representation and the disaggregation algorithm. In the human-visualization interaction side, the *filter* and *exploration* interaction techniques enable the user to configure views according to his mental map (conditioned by its previous knowledge), to correlate context variables in large buildings (e.g., temperature, variables from important subsystems or occupancy) and to evaluate the obtained decompositions. For instance, the correlation between an individual consumption and the ambient temperature is insightfully assessed if users are able to filter the represented data by ranges of temperature. In addition to the filter and exploration interactions, low-latency view transitions are also desirable in order to provide a fluid user-representation loop [126].

The filter and exploration mechanisms intrinsically entail iterating all data items. In scenarios with a large amount of samples, iterating all data items may cause long latencies, and therefore it makes fluid interactions difficult. This introduces the requirement of an efficient indexing of the data capable of speeding up the filter and exploration mechanisms.

At this regard, *Data Cubes* (OLAP, *on-line analytical processing*) [144, 145] have demonstrated to be efficient data structures for storing multidimensional *records* from table-based datasets. Data cube structures are commonly explained with a specific terminology that was firstly introduced in [144, 145]. In [146], this initial terminology was reviewed and adapted to energy monitoring problems. We will adopt the terminology in [146] for our integration of the energy disaggregation into the VA paradigm.

The table-based data shown in Figure 3.16 will be used to explain the elements and operations of data cube structures. As it can be seen in the figure, the NMF outcome is included in the table as new columns of the original table with the context factors of the building (e.g., external temperature or information from critical subsystems of the building).

The individual consumptions obtained after flattening the product $\mathbf{w}_i \mathbf{h}_i^T$, the basis consumptions \mathbf{w}_i and their activations \mathbf{h}_i^T from the NMF-based NILM analysis are examples of these added columns. Note that $\mathbf{w}_i \in \mathbb{R}^{M \times 1}$ and $\mathbf{h}_i^T \in \mathbb{R}^{1 \times N}$ do not have the same elements as the $M \times N$ number of records in the table, so that they should be expanded. Thus, \mathbf{w}_i (daily basis) is cyclically repeated N times, while each element of \mathbf{h}_i^T (daily activations) is repeated M times in order to naturally fit both vectors in the table.

Taking into account this table as the starting point, the data cube elements are defined as follows:

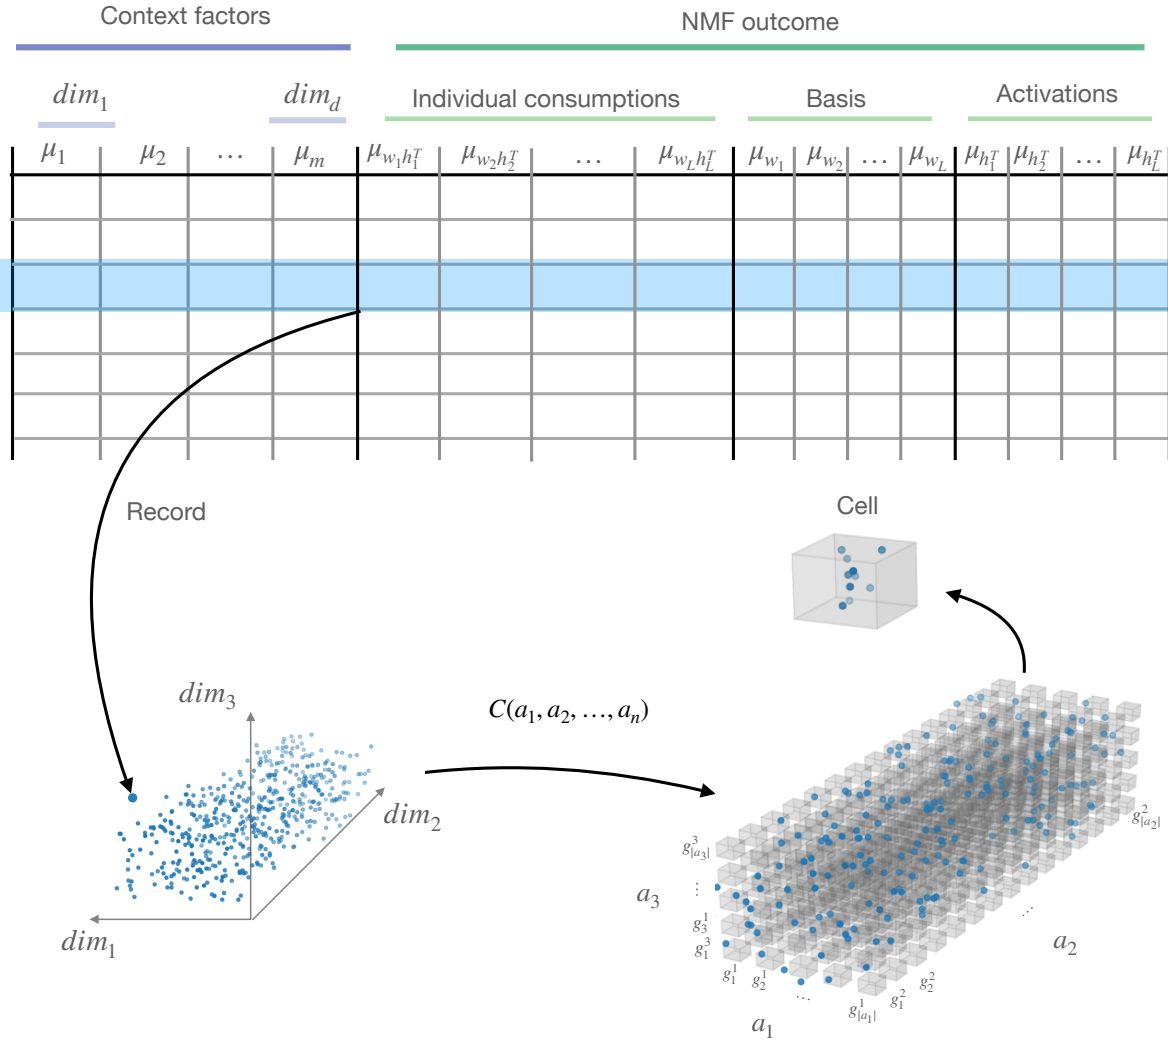


Figure 3.16: Intuition of data cube formation and terminology.

- *Measures* are defined as the numeric measurements $\mu_1, \mu_2, \dots, \mu_m$ gathered from the system. The main active power energy demand in kW , the active power from the individual consumptions or the ambient temperature are examples of measures. The measures are samples from *dimensions* that provide semantics to the scalar values registered in the table. Hence, for instance, the main active power energy demand and the individual active power measures belong to the same dimension *active power (kW)*.

- *Attributes* are the result of applying a *grouping operation* to the measures, so that their samples are rearranged into a finite number of *groups* or subsets. In the case of continuous dimensions, measures can be grouped on contiguous bins of a specific range (e.g., grouping ambient temperature in bins of 10°C degrees each). By contrast, some dimensions can be grouped exploiting their hierarchical nature. For example, time-based dimensions can be naturally grouped by minutes, days or months. After applying the grouping operation, attributes can be formally defined as a set of groups as follows:

$$a_i = \{g_1^i, g_2^i, \dots, g_{|a_i|}^i\}$$

where g_j^i denotes the j -th group of the i -th dimension and $|a_i|$ denotes the cardinality of the attribute a_i . Each record belongs to one and only one of the groups of the attribute.

- *Data cube* is the resulting data structure after arranging the samples of the initial table in cells, whose coordinates are the groups of the defined attributes. Hence, the data cube is made of $|a_1| \times |a_2| \times \dots \times |a_n|$ cells, and it is denoted as:

$$C(a_1, a_2, \dots, a_n)$$

In the resulting data cube C , each sample of the initial table is mapped to a cell whose coordinates are (g^1, g^2, \dots, g^k) , being g^k the group to which the sample belongs for the k -th attribute. The cells form an hypercube, whose dimensions (do not confuse with the dimensions of the measurements) are defined by the attributes, as it is shown in Figure 3.16. Note that a cell can host more than one sample, but it can also be empty.

Once the data structure is formed, a set of algebraic operations can be applied to the hypercube to adjust the output values of the cube to the user's intentions. Hence, the stored data in the cube can be prepared for the aforementioned exploration, filtering and selection mechanisms by means of the following operations:

- *Aggregation*. After the cube creation, an output information for each cell should be provided, which can be either directly the array of the values placed in the cell (see

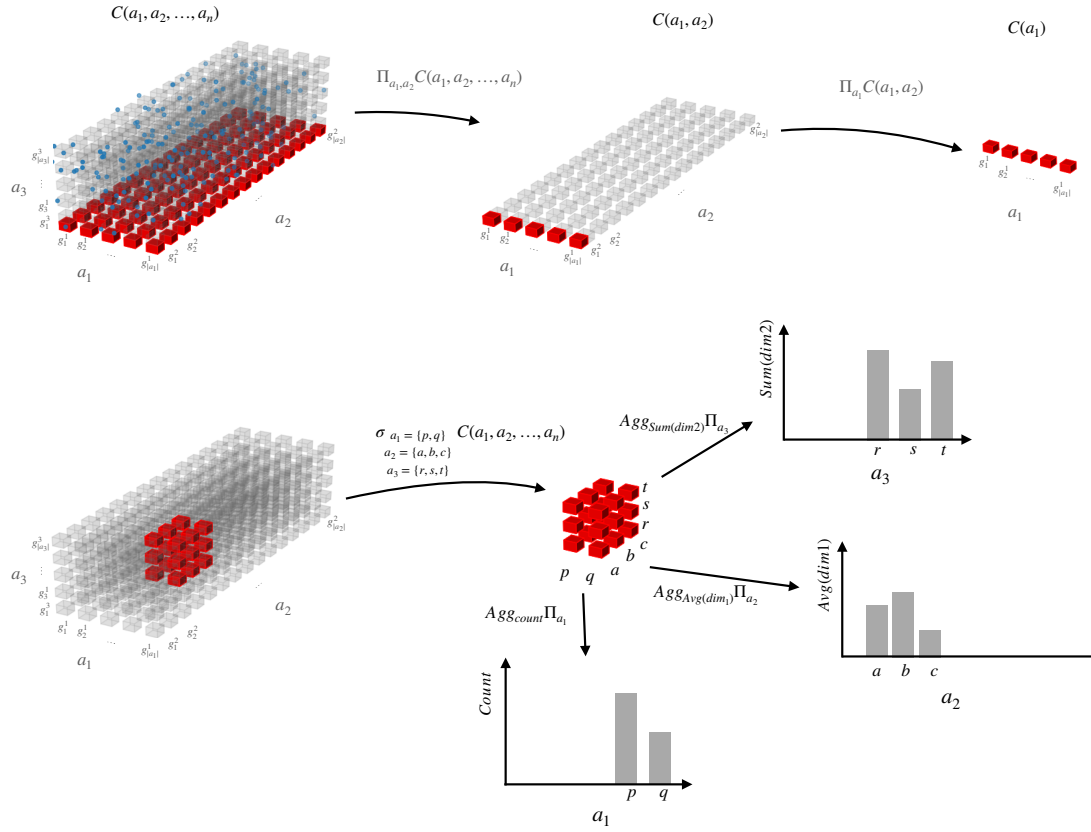


Figure 3.17: Examples of data cube operations. In the first row the *projection operation* is shown by firstly projecting the initial cube to a 2D cube and then to 1D cube. In the second row a combination of *selection*, *projection* and *aggregation* operations is shown.

Figure 3.16), or a representative value that summarizes them. The aggregation operation computes these cell-representative values, and it is formally defined as:

$$\text{Agg}_{\mathcal{A}} C(a_1, a_2, \dots, a_n)$$

where \mathcal{A} is an aggregation function to be computed over the cell values, such as count, sum or mean, and it is considered a user-defined parameter. In addition to the \mathcal{A} function, the user can also decide on which dimension the aggregation is computed (note that for count aggregations, this is not needed). Thus, the average aggregation by the dimension dim1 would be $\text{Agg}_{\text{avg}(\text{dim1})}$. The aggregation operation returns a n -array with $|a_1| \times |a_2| \times \dots \times |a_n|$ elements.

- *Projection* operation enables the user to select the attributes used to analyze the cube. This allows users to define different perspectives of the data, flattening the

initial cube into another cube with less number of cells. In [146], the projection operation was defined as:

$$\Pi_{a_1, a_2, \dots, a_p} C(a_1, a_2, \dots, a_n) \rightarrow C(a_1, a_2, \dots, a_p)$$

where $a_1, a_2, \dots, a_p \in a_1, a_2, \dots, a_n$ is a subset of the original attributes. In Figure 3.17, the projection operation is illustrated by two examples. Firstly, the initial cube is squashed into two dimensions, and then into one. After projecting, all the records are reorganized into a new reduced set of cells, causing an increment of the records stored in each cells. The projection operation is related to the exploration mechanisms, since it enables the user to set a subproblem (perspective) where only a few attributes are involved. As in the case of the aggregation, the set of attributes a_1, a_2, \dots, a_p are interactively defined by the user.

- *Selection* operation filters the records according to a user-defined logical condition. Typically, this condition is the selection of a subset of groups in one or several attributes $a_i = \{a, b, c\}, a_l = \{p, q\}$. More formally, the selection operation can be expressed as:

$$\sigma_{\substack{a_l = \{p, q\} \\ a_i = \{a, b, c\}}} C(a_1, a_2, \dots, a_n) \rightarrow C(a_1, a_2, \dots, \underbrace{\{a, b, c\}}_{a_i}, \dots, \underbrace{\{p, q\}}_{a_l}, \dots, a_n)$$

Particularly, if a contiguous subset of groups is selected, a *dice* is extracted from the initial cube, as it is shown in Figure 3.17. The contiguous selections are useful in continuous dimensions, where the analysis of contiguous elements is coherent. The selection operation also enables non-contiguous filtering which may be more suitable for categorical data, where contiguity is less relevant. The *Selection operation* is tightly related to the *filter interaction* pathways, since users can filter the data to be represented according to their intentions.

All these operations can be combined to produce very particular perspectives, such as the cube perspective shown in the second row of Figure 3.17:

$$\text{Agg}_{\text{avg}(\text{dim}_1)} \Pi_{a_2} \sigma_{\substack{a_1 = \{p, q\} \\ a_2 = \{a, b, c\} \\ a_3 = \{r, s, t\}}} C(a_1, a_2, \dots, a_n) \quad (3.10)$$

where a small dice is initially extracted from the cube by selecting (filtering) a few attributes, for then projecting the records of the dice into the a_2 attribute. After projecting, the output cell-representative values, represented as barcharts in Figure 3.17, are computed by averaging the dim_1 values of each cell records.

Cell indexing reduces the computational time in the selection and projection operations since, for instance, the infinite possible measurements of a continuous dimension are reduced to a finite number of groups, when it is grouped into an attribute. Therefore, the number of items to be iterated in a selection process is notably reduced. This *binned aggregation principle* [147] enables to operate with data cubes created from large datasets (more than a million of records) and to visualize the outcome of the cube on “live” representations [146, 148, 149]. Note, for instance, that all data cube operations (aggregation, projection, selection) involved in (3.10) could be *steered* by the user in real time during the analysis.

Integrating both raw demand data and ML processed information (NMF) in the same cube, we expect an increment in the interpretation of the disaggregation, since insightful correlations between the resulting NMF latent patterns and the context information can be found in a highly interactive way through dice and projection operations.

3.3.4 Data visualization techniques

Although the data cube approach provides the user with efficient interaction mechanisms with the data, the computed outcome of the cube is composed of numeric values, which have to be visually encoded according to the principles mentioned in section 3.3.1 in order to produce a view I , following our problem formulation. In [148], the authors suggest specific visual encodings suitable for representing aggregated values gathered from similar data cube projections. Examples of these data visualization techniques are *barcharts*, *heatmaps* or *scatter plots*.

In [146], the authors represent the data cube outcome with a set of barcharts in a small multiples layout, providing interactive elements in the dashboard (e.g., brushing tools, text boxes, etc.) that allow the user to change the perspective of the data by triggering data cube operations.

In our approach, the combination of small multiples, interactive mechanisms and data cube operations in [146] is extended. Thus, we integrate new 2D projections of

the cube in order to provide a bidimensional representation of data. For instance, 2D projections of time variables, such as the pairs of attributes DayOfWeek/Hours or DayOfWeek/WeekOfYear, provide insightful time views in the form of calendars of weekly views. Moreover, these 2D time views can be easily connected to the NMF attributes by means of the selection and aggregation operations.

Taking this into account, the following data visualization techniques were suggested to visually encode the data cube outcome:

- *Barchart*. The so-called *barchart* technique, which is shown in Figure 3.18, is used for representing 1D projections of the data cube, in which the resulting cell-aggregated values are represented by means of the efficient mark-channel combination formed by the height (channel) and the bars (mark) (see Figure 3.14). At the top of each barchart in Figure 3.18, there are three combo boxes by which the user defines both the projection Π and aggregation operations $\text{Agg}_{\mathcal{A}(dim)}$. The first combo box defines the attribute on which the data cube is projected, whereas the second combo box indicates which aggregation function is applied. If the selected aggregation function \mathcal{A} is either mean or sum, the dimension to be aggregated dim is defined by the third combo box. The selection operation σ is implemented by means of *brush* and *drag* gestures, as it is shown in Figure 3.18. The brush event defines a selection of contiguous values (dice selection) in the attribute represented in the x -axis. This selection can be quickly redefined by means of drag events, so that the user can filter the analysis to different regions of interest, being able to spot insightful visual correlations in the rest of the represented projections.
- *Heatmap*. An advisable visualization for representing the aforementioned 2D projections are the binned heatmaps [148], as shown in Figure 3.19. In this kind of visualizations, each value of the outcome of a 2D cube projection is visually represented by a square, whose color encodes the resulting aggregated value. Suggested heatmaps integrate two combo boxes whereby the user only defines the aggregation operation. The projection and selection operations are disabled, since the projection attributes are fixed to a specific pair of attributes and the selection operation is delegated to brush gestures in other barchart views.

Thus, we suggest two heatmap views: 1) DayOfWeek/Hour projection (week view), shown in Figure 3.19a; and 2) a DayOfWeek/WeekOfYear projection (calendar view), shown in Figure 3.19b. Both views are widely used visualizations and fit perfectly

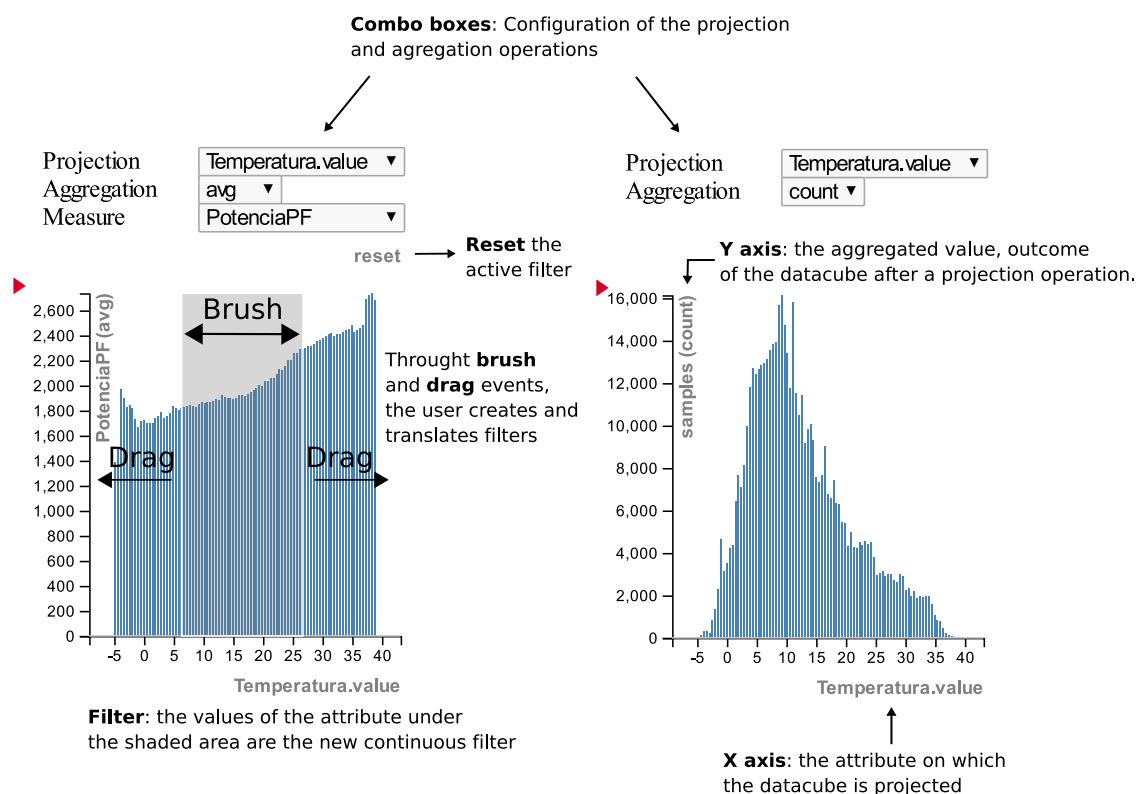


Figure 3.18: Interactive barchart technique for representing 1D projections of the data cube.

in the energy load monitoring for large buildings due to the strong seasonality of their data. The calendar and weekly layouts provide insight on many kinds of periodicities closely related to the human activity that takes place in large facilities.

- *Sparklines*. In order to improve the interpretation of the NMF outcome, we suggest representing the basis consumptions w_i , integrating the sparklines visualization shown Figure 3.20. Representing the basis consumptions through sparklines allows the final user to spot and compare the learned daily patterns, using a small area of the dashboard, following the principle of small multiples and the maximization of data density in the graphic.

Applying these data visualization approaches to the NMF + cube paradigm, the user can visually explore the NMF decomposition together with other context factors recorded from the building, being able to get useful correlations between them. This results in a significant cross-fertilization between both approaches that leads users to gain a knowledge that cannot be extracted from NMF analysis or data visualization

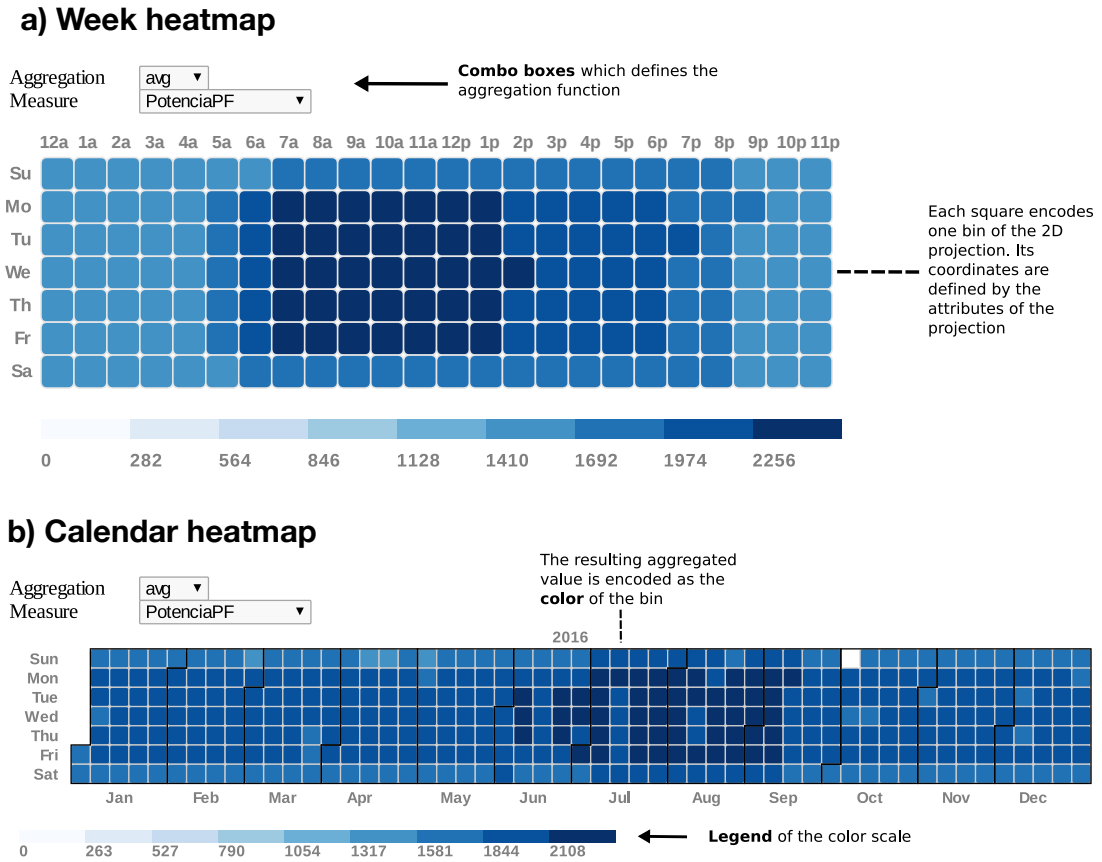


Figure 3.19: Heatmaps which represent the 2D projections of the data cube. a) represents the average of the total power demand of the hospital (PotenciaPF) grouped by hour and day of week. b) represents the average of the total power demand grouped by days of year arranged in a calendar distribution.

techniques, separately.

3.4 Experiments

The NMF-based energy disaggregation, its visualization and the integration of both approaches in a VA application by means of the interaction mechanisms provided by the data cube were tested with real electric data from the Hospital of Leon dataset described in section 3.1.1. In the remainder of the section, we will discuss the NMF decomposition of the Hospital main energy demand according to different number of components L . Then, we shall present the most relevant patterns discovered through NMF and their connections with relevant events in the Hospital. Finally, we will show how powerful the VA integration is in order to interpret the NMF decomposition by means of several use

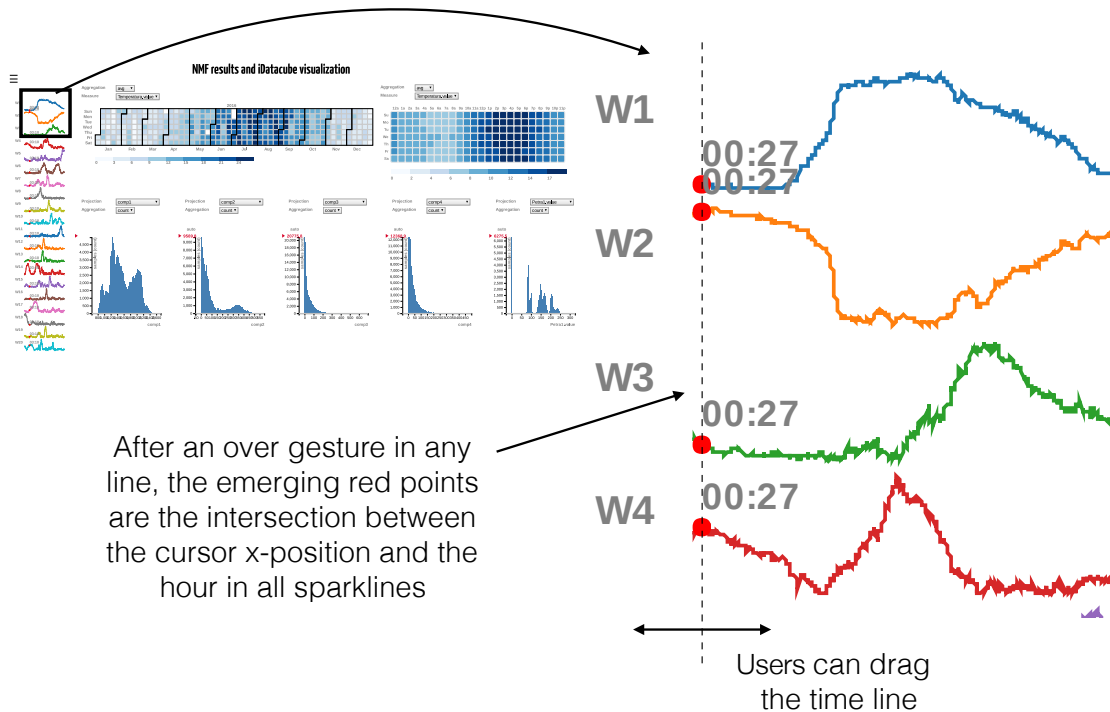


Figure 3.20: Snapshot of sparklines of NMF basis consumptions.

cases.

All the results and most of the figures of this section have been published in the journal *Energy and Buildings* within the article [10].

3.4.1 Experimental set-up

Subsequent experiments in this section are based on analyzing the NMF decomposition applied to the main active power from the Hospital of Leon dataset. As it was previously mentioned, the active power and the subsequent NMF decomposition are not the only data introduced in the data cube, but also other *context factors* are included in the creation of the cube. Among the vast amount of context factors measured in the hospital, we categorized them in three different types: 1) Time-related variables; 2) Temperatures; and 3) Variables from critical subsystems of the hospital.

Hence, time variables such as DayOfYear or Hour are included, together with the ambient Temperature of the hospital. Regarding the critical subsystems, we will focus on variables from the heating, ventilation and air conditioning systems (HVAC) made of seven cooling machines, which feed with cold water the ring that keeps the diagnostic devices chilled. The HVAC subsystem is essential for the normal functioning of the hospital, and it represents about 30% of the overall consumption of the facility.

Hence, an increment of the knowledge about its behavior could be a valuable feedback that may encourage energy efficiency improvements through changes in the hospital configuration, such as HVAC parameterization or more efficient planning of the hospital's activity. For these reasons, energy consumptions of the cooling machines are included, so that the variables under study in our experiments are shown in Table 3.4.

Table 3.4: Variables used as *context factors* in visual analytics approach. The outcome of the NMF decomposition will be added to these variables before creating the data cube.

Label	Variable description	Units
Acometida	Main active power	kW
P1	Cooling machine 1	kW
P2	Cooling machine 2	kW
P3	Cooling machine 3	kW
P4	Cooling machine 4	kW
P5	Cooling machine 5	kW
T1	Cooling machine 6	kW
T2	Cooling machine 7	kW
Temperature	Temperature	°C
DayOfYear	Day of year	No. of the day (0, 1, ..., 364)
Hour	Hour	No. of the hours (0, 1, ..., 23)
Month	Month	No. of month (0, 1, ..., 11)
WeekOfYear	Week of year	No. of week (0, 1, 2, ..., 11)
DayOfWeek	Day of week	No. of day in week (0, 1, 2, ..., 6)

Using the data described in Table 3.4, the NMF decomposition is computed and arranged in a data cube. Once the data and the NMF decomposition are integrated in a data cube, the visual energy monitoring tool shown in Figure 3.21 is developed, using the aforementioned visual encoding principles and visual techniques.

Regarding the NMF computation, only the main active power (Acometida) from Table 3.4 was the input to the NMF model. The remaining variables were only incorporated into the data cube structure as context variables, from which users can operate and produce

new perspectives. As it was mentioned in section 3.1.1, the data from the Hospital of León were recorded with a sample period of one minute, so that the obtained dataset contains a huge amount of records which might cause long latencies in the initial computation of the NMF and the data cube. In order to speed up the initial load of the application, we suggested downsampling the records with a 1:3 downsample operation (average as aggregation function), reducing the resolution of the data to three minutes.

Applying NMF decomposition requires to split the time series *Acometida* into windows in order to generate the columns of the input matrix \mathbf{V} described in (3.8). We set the size of the windows so that the obtained columns were daily observations of the total demand of the hospital. With this arrangement, the resulting coefficient vector \mathbf{h}_i represents the daily contributions of the basis consumption \mathbf{w}_i throughout the year.

In the NMF decomposition, we discarded those days which presented missing records, because the NMF algorithm would learn the gaps in the data as components, instead of other relevant patterns. Both the preprocessing of the data and the NMF decomposition were developed in a *Python* script which makes use of several well-known data analysis libraries such as *Pandas* [150] or *Scikit-learn* [151].

After the NMF computation, the context factors in Table 3.4, together with the NMF decomposition are arranged in a data cube. The data cube creation and its operations are implemented with the *Javascript* library *crossfilter* [152]. Once the data cube is generated, its output values are visualized using the data visualization techniques of section 3.3.4, which were also implemented in *Javascript*, using the well-known data visualization library *D3.js*. As a result, the dashboard shown in Figure 3.21 is presented to the users in a web page interface. This web page is hosted in the python web server *Flask* [153], that enables us to integrate the python implementation of NMF with the data cube and visualization implemented in Javascript.

The hyperparameters of the NMF algorithm, such as the number of components L or the maximum number of iterations, are chosen by the user in a dialog window, as it is shown in Figure 3.22. By default, we suggest the configuration shown in Figure 3.22, which ensures a representative set of sparse components.

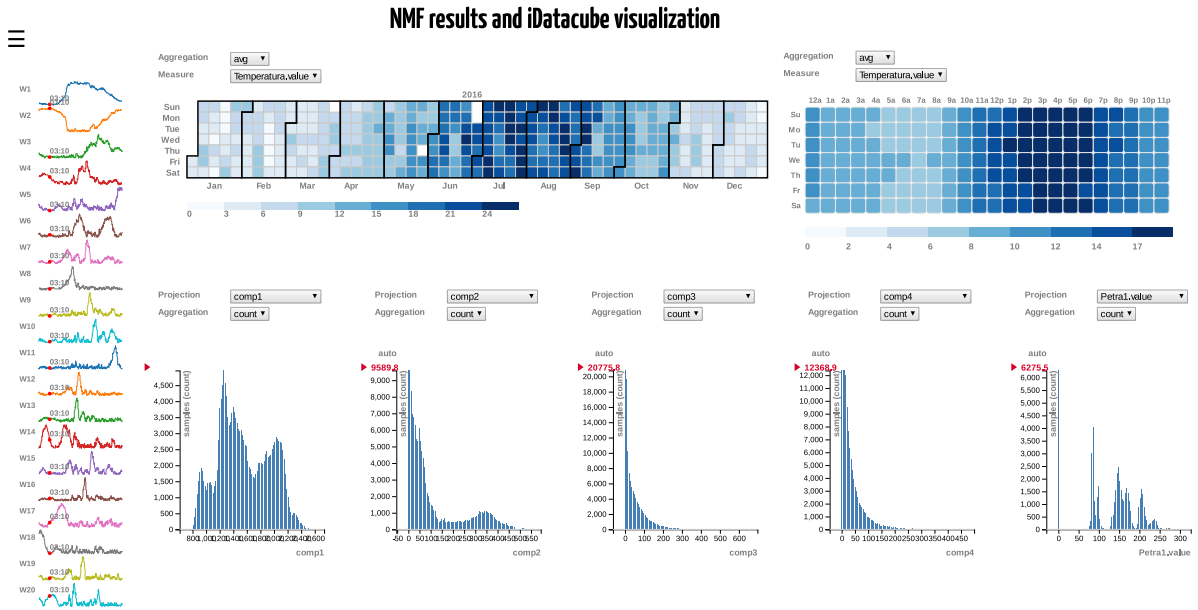


Figure 3.21: General view of the suggested VA application.

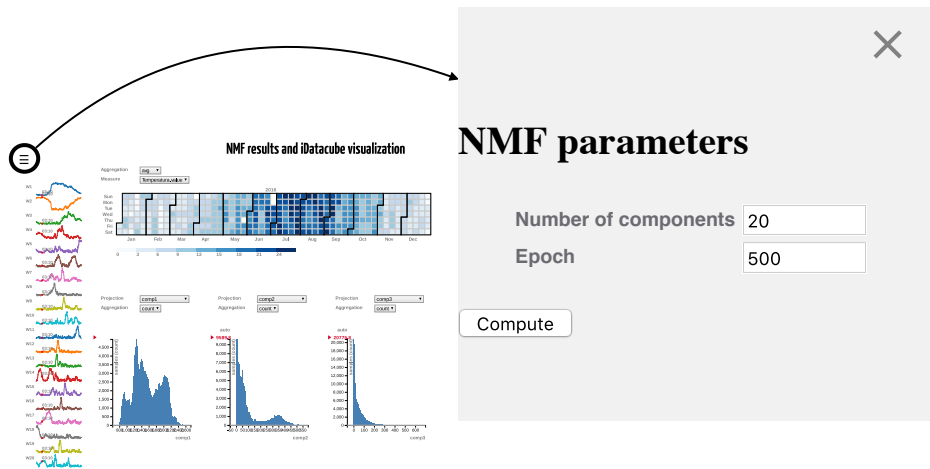


Figure 3.22: Dialog window which enables the user to interactively choose the parameters of the NMF decomposition.

3.4.2 Analysis and interpretation of NMF components

Firstly, we analyze the NMF decomposition alone, by training several NMF models with different number of components L . After the computation of all the NMF models, the resulting decompositions are evaluated by means of the *Frobenius norm* of $\mathbf{V} - \mathbf{WH}$ as the reconstruction error:




Figure 3.23: Reconstruction error according to the number of components.

$$\text{reconstruction error} = \|\mathbf{V} - \mathbf{WH}\|_F^2 \quad (3.11)$$

A total of 29 models with $L = 2, 3, \dots, 30$ were trained, and, in Figure 3.23, the reconstruction errors for all models are shown. Reconstruction errors sharply decrease from $L = 1$ to $L = 20$, and for values larger than $L = 20$, the error decreases slowly. Taking into account that the larger the number of components, the lesser the interpretability of the decomposition, we consider that a decomposition in 20 components is accurate enough. Therefore, L is set to 20 by default from now on.

Once the L parameter is set, the obtained basis consumptions and their activations are evaluated in Figure 3.24. The computed basis consumptions are sorted by order of importance, so that the first bases, such as \mathbf{w}_1 , \mathbf{w}_2 and \mathbf{w}_3 , are the most relevant ones. Analyzing their corresponding activations, it is observed that they are associated with weekly patterns in the case of \mathbf{h}_1 and \mathbf{h}_2 , and seasonal patterns in the case of \mathbf{h}_3 .

From the fourth component onwards, the bases are related to improbable events that show sparse bases and activations (short consumptions relevant only in specific days), and therefore they can be attached to abnormal behaviors or occasional faults.

Taking a closer look at the most relevant factor ($\mathbf{w}_1, \mathbf{h}_1$), one can observe that its activations \mathbf{h}_1 extend over all days of the year, but specially in workdays. Regarding

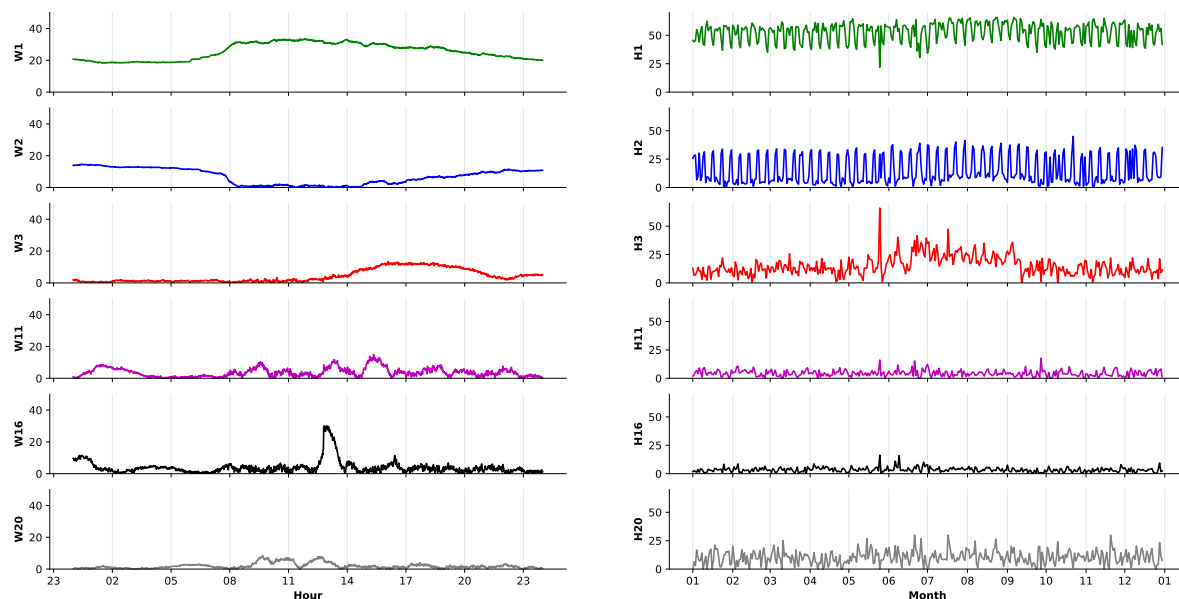


Figure 3.24: Some components from the trained NMF model with $L = 20$. The first three factors are more relevant both in daily profile and throughout the year than the last factors.

its basis consumption \mathbf{w}_1 , it shows a hump shape similar to a daily profile of energy demand of the hospital. Hence, it can be considered as a “baseline” on which the rest of the components will be aggregated.

Complementing the aforementioned first factor, the second factor ($\mathbf{w}_2, \mathbf{h}_2$) is only relevant during the weekends and holidays. Regarding its basis consumption \mathbf{w}_2 , it is also complementary to \mathbf{w}_1 , reaching its maximum values from 5 p.m. to 8 a.m. (inactive hours of the day). This component, when is summed to the first component, makes the reconstructed daily profile smoother, reducing the hump shape. By means of the second component, the NMF decomposition is actually classifying the daily profiles into workdays where the activity of the hospital raises during the central hours of the day, and weekend days where the hospital activity remains more constant and the hump shape is less predominant.

The third factor ($\mathbf{w}_3, \mathbf{h}_3$) can be associated with the seasonality of the hospital, since its activations \mathbf{h}_3 increase during the summer. Its daily basis \mathbf{w}_3 presents large values of consumption from 2:00 p.m. to 8:30 p.m., which reveals that the hospital draws more energy during the summer afternoons. Additionally, in \mathbf{h}_3 a large peak is observed at the end of May, which coincides with a sharp decrease of \mathbf{h}_1 . The NMF decomposition

highlights this day as an abnormal profile which had a large consumption in the afternoon and evening.

In the last factors, the magnitude of the activations decrease, reducing their importance in the final reconstruction. Hence, from the tenth factor onwards the basis consumptions are only relevant in specific days of the year. For example, the activations \mathbf{h}_{11} , \mathbf{h}_{16} and \mathbf{h}_{20} , whose basis consumptions show peaks for a few hours and lack of any periodicity. This sparsity in the basis consumptions and their activations is related to the natural sparseness of NMF decomposition. The interpretation of these sparse factors, from the energy demand point of view, is related to non-periodic patterns that may be caused by malfunctions or novelties.

In the light of these results, we can conclude that the NMF decomposition by itself is capable of learning periodic and non-periodic patterns and localizing them in time, but it is still difficult to associate them with specific individual consumption from the facility or to recognizable events. At this regard, the integration of the expert knowledge by means of the data cube and the data visualization can substantially improve the NMF disaggregation analysis.

3.4.3 Evaluation of the visual analytics application through use cases

In order to assess the performance of the human-in-loop analysis provided by our VA approach of 3.21, in this section several use cases will be presented. These use cases show that NMF factors can be insightfully contextualized by interacting through the visualization parameters S_V and the parameters related to the NMF decomposition S_D . Within the suggested dashboard, the user can tweak the following S_V and S_D parameters:

- S_D . The user can modify the following NMF-related parameters:
 - Number of factors L . Although it is set to 20 by default, it can be redefined as shown in Figure 3.22.
 - Number of iterations in the NMF optimization.
- S_V . The user can change the view by means of:

- Projection operations Π triggered by the combo boxes in each of the interactive barcharts.
- Aggregation operations Agg defined by the aggregation function and the dimension on which the aggregation function is applied. Both the aggregation function and the dimension are also defined in the barchart and heatmaps combo boxes.
- Filter operations σ triggered by *drag* and *brush* gestures on barcharts.

In the following use cases we will refer to these parameters to explain the pathways of interaction followed by the user to configure the explored views.

Use case 1: filtering by large values of the components

A natural question for understanding any NMF component is *when is this component relevant?* or *what occurs in the facility when the component is high?* These questions can be answered by filtering those records that have the highest NMF individual component values $\mathbf{w}_i \mathbf{h}_i^T$ (in the application we refer to these individual component as *comp*). In other words, the analysis is focused on the records where the component is highly relevant, with the aim of finding recognizable patterns in other variables, such as the time context factors (*when is it relevant?*).

In Figure 3.25, three examples of filtering by large values of the first three components are shown. The filtering operation is performed by brushing on a barchart previously configured as a projection of the component under analysis. In Figure 3.25a, after filtering by large values of the first component, a weekday pattern emerges in the calendar heatmap using $\text{Agg}_{\text{count}}$ as aggregation function. In other words, the calendar indicates which days host the largest values of the component. This corroborates, in an intuitive way, the aforementioned analysis of Figure 3.24.

In Figure 3.25b, the same filter operation is applied, but this time over the second component. As expected, a weekend pattern emerges in both week and calendar heatmaps, using configured $\text{Agg}_{\text{count}}$ as the aggregation operation, as well. Although the highlighted information is the same as the shown in Figure 3.24, one can identify the holiday patterns in a familiar visualization, as the calendar heatmap, easier than in a static view as Figure 3.24. Finally, in Figure 3.25c, seasonal patterns during afternoons and evenings emerge after filtering the largest values of the third component.

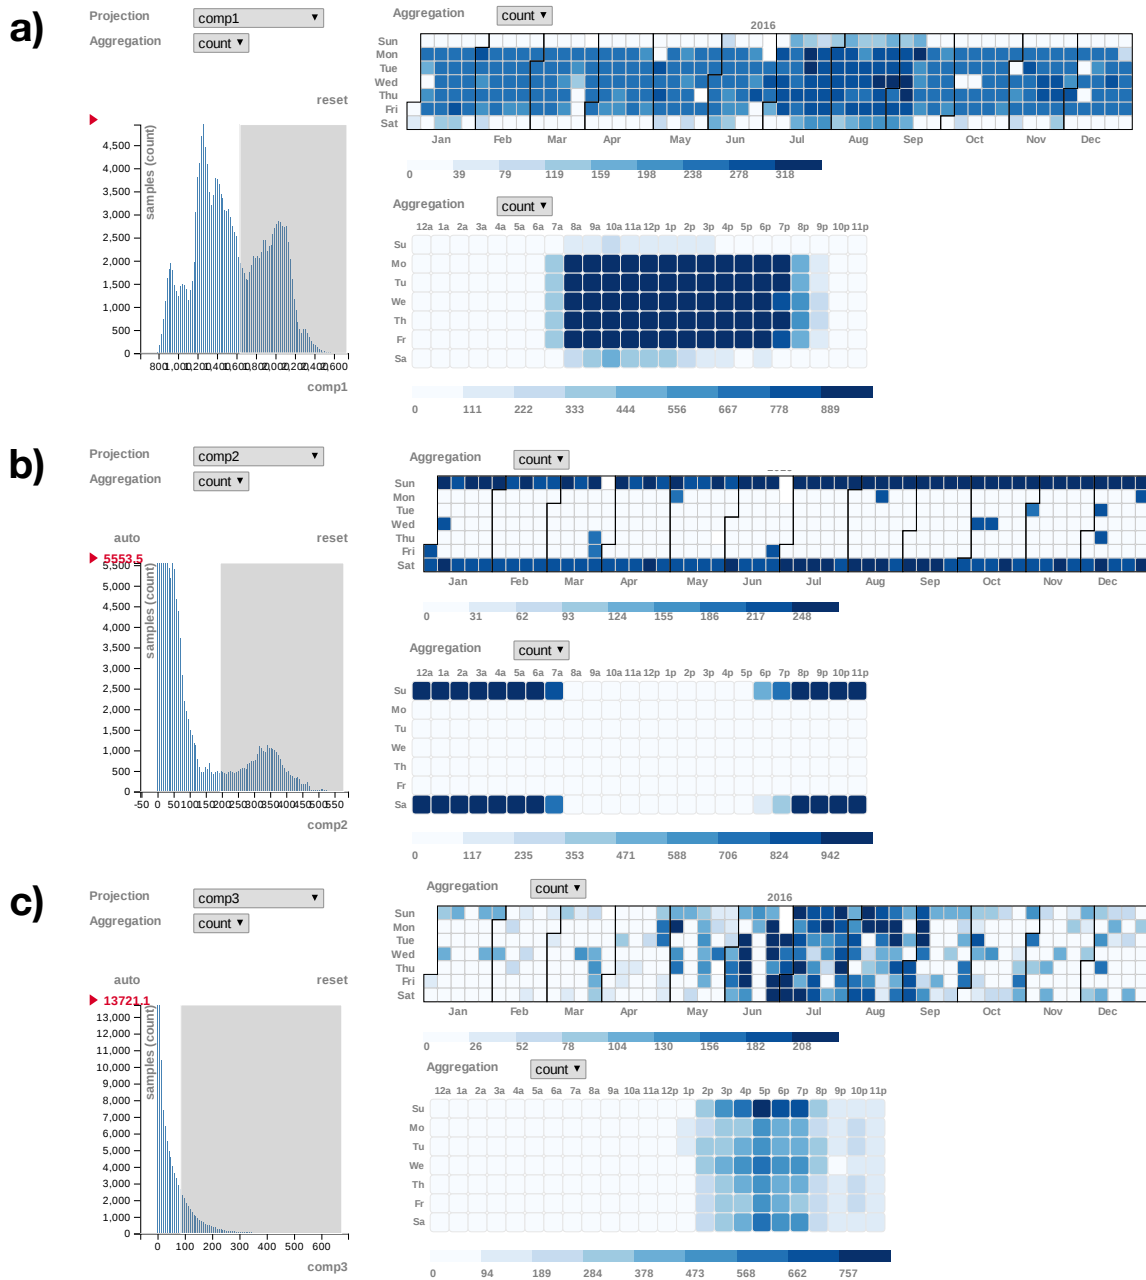


Figure 3.25: Filtering the largest values of the first three components by means of drag gestures in their associated barcharts. a) Weekday patterns are spotted in both the calendar and weekly heatmaps after filtering by the largest values of the first component and selecting count as the aggregation function. In b), the same analysis as in a) shows that the highest values of the second component highlight the weekends and holidays. The calendar in c) shows that the largest values of the third component are related to a seasonal behavior.

Despite the findings are not very surprising, filtering by large values of the components is a convenient first step in an overview-to-detail analysis, since these findings may raise new questions or hypotheses, leading to a new cycle of interaction in our human-in-loop analysis. For instance, the filtering operation on the third component lead us to new questions, such as: “*why is the third component higher in the afternoons?*” or “*does this component have any relation to the cooling machines?*”.

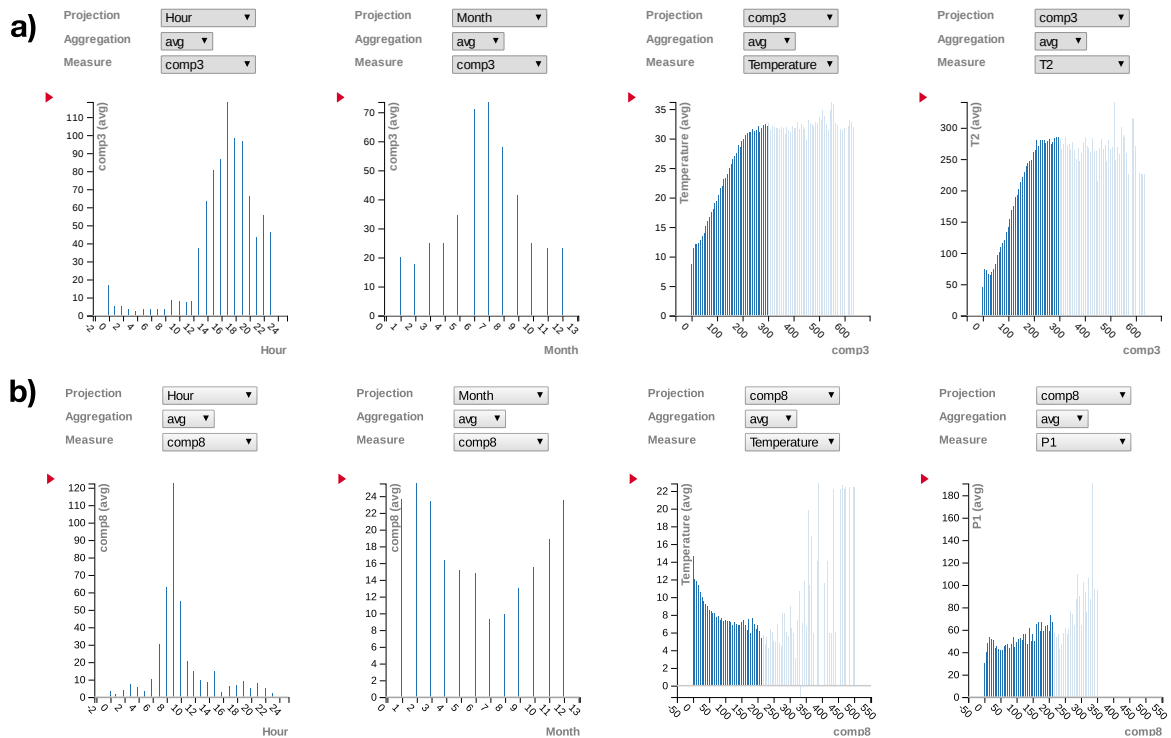


Figure 3.26: Configuration of the interactive data cube whereby an insightful visual correlation was found in the third and eighth components. In a), correlations can be established between the third component, the ambient temperature and the consumptions of the cooling machines of type T. In the same way, b) shows how the low and medium values of the 8-th component can be strongly related to low temperatures and slightly related to the consumption of the P1 machine.

Use case 2: visual correlation between components and the rest of the attributes

The combination of data visualization techniques and data cube projection/aggregation operations enables the user to graphically establish correlations between attributes. Although these correlations may not imply causation, they reveal hidden relations between the context factors and the computed NMF components, that an expert eye can easily connect to events or specific behaviors of the facility.

These graphical correlations emerge by configuring a barchart to represent the data cube operation $\text{Agg}_{\text{avg}(\text{dim1})}\Pi_{a2}C$, where the cube is projected into the attribute $a2$ and the aggregation function returns the averaged values computed over the dimension dim1 , different from the dimension of the projection attribute dim2 . Thus, if the two dimensions are correlated, the operation $\text{Agg}_{\text{avg}(\text{dim1})}\Pi_{a2}C$ returns a barchart, where the y -axis values monotonically increase or decrease along the x -axis.

Following this procedure, the correlation of the NMF components with the context factors can be evaluated in the different barchart views of our dashboard. In Figure 3.26, the correlation between 3-th and 8-th components with other attributes are analyzed. Firstly, the correlations between the third component, the months and the hours are analyzed in the first two barcharts of Figure 3.26a. These correlations are obtained by projecting on the attributes hour and month and aggregating by the third component.

The resulting barcharts show a strong influence of the third component in the summer afternoons, which suggests that the component may be related to the ambient temperature. In order to clarify this hypothesis, in the following barchart the correlation between the third component and the ambient temperature is obtained from the operation $\text{Agg}_{\text{avg}(\text{Temperature})}\Pi_{\text{comp3}}C$, revealing a strong local correlation between both attributes.

This finding led us to a new hypothesis *Is the third component related to some cooling machine?*. This question is answered in the last barchart of Figure 3.26a by means of the operation $\text{Agg}_{\text{avg}(\text{T2})}\Pi_{\text{comp3}}C$, that shows a significant positive and local correlation between the consumption of the cooling machine T2 and the third component.

Analogously, in Figure 3.26b, one can observe that the 8-th component shows a peak of consumption at 9:00 a.m. during winter months. Following the same analysis as in the 3-th component, we observe that the component is inversely correlated to the Temperature and slightly correlated to the energy demand of the cooling machine P1. Bearing these findings in mind, the morning peak of consumption highlighted by the 8-th component seems to be correlated to the start-up of P1.

Graphical correlations shown in Figure 3.26 differ from analytically extracted metrics of correlation (e.g., *Pearson coefficient*) in that they provide the user with a fine-grained analysis instead of a unique value, so that the correlation can be locally established by ranges of values. For instance, in Figure 3.26a, the correlation between the third

component and T2 is stronger for low and medium values of the component, and weaker for high values of the component.

The proposed visual-based correlation analysis for the 3-th and 8-th components is assessed in Tables 3.5 and 3.6, where the Pearson coefficients of the components against the consumptions of the machines P1, T2 and the ambient temperature are computed. The resulting Pearson coefficients are conditioned to the following ranges of values of the components: low values (L) between [0-100], medium values (M) between [100-200], high values (H) between [200-350] and the whole range (WR) of values [0-350]. Pearson coefficients in Table 3.6 are computed from the output values of the data cube (after the grouping, projection and aggregation operations), while the coefficients of the Table 3.5 are directly computed from the raw samples of the initial table, before generating the data cube.

Table 3.6 evaluates the visual-based correlations and 3.5 the sample-based correlations, showing that both correlations have the same sign, and the visual-based correlations are stronger than the sample-based ones. This may be explained by the smooth effect of averaging the cell records, which attenuates the effect of noise and local variations. This comparison confirms the reliability of visual correlations, since they are close to those calculated analytically.

A limitation of the visual-based correlation methods emerges when the cells do not host a representative number of records. This leads to non-representative aggregations that result in visual noise in the barcharts. An example of this noise can be seen in the highest values of the components in Figures 3.26a and 3.26b, where the aggregated values (represented by the height of the bars) show odd spikes. Bearing this mind, the data cube cells that host less than 100 records are then represented by bars with high transparency.

The former correlation use cases exemplify the human-in-loop paradigm, where several changes of view encouraged by new hypotheses take place and, eventually, a relevant and reliable information is extracted. For instance, in the previous analysis, it was found that, in cold days the hospital presents a peak of consumption at 9:00 a.m. highlighted by the 8-th component that is weakly related to the energy demand of P1 machine, while the consumption of the machine T2 is related to the increase of the total consumption on summer afternoons, as indicated in 3-th component.

	comp3				comp8				
	L	M	H	WR	L	M	H	WR	
P1	0.03	-0.10	-0.04	-0.10	0.09	0.05	0.08	0.09	1
T2	0.22	0.27	0.03	0.46	-0.15	-0.08	0.02	-0.15	0
Temp	0.41	0.32	0.13	0.63	-0.21	-0.04	0.04	-0.22	-1

Table 3.5: Pearson coefficients computed from the raw samples of the initial table, and between the 3-th, 8-th components, the individual consumptions of the cooling machines P1 and T2, and the ambient temperature conditioned to low (L) [100-200], medium (M) [100-200], high (H) [100-340] values of the component, as well as the whole-range correlation (WR) [0-350].

	comp3				comp8				
	L	M	H	WR	L	M	H	WR	
P1	-0.38	-0.96	-0.60	-0.91	0.25	0.70	0.39	0.61	1
T2	0.95	0.99	0.03	0.94	-0.84	-0.78	0.48	-0.44	0
Temp	0.98	0.99	0.75	0.96	-0.93	-0.48	0.48	-0.48	-1

Table 3.6: Group-based Pearson coefficients computed from the data cube output, and between the 3-th, 8-th components, the individual consumptions of the cooling machines P1 and T2, and the ambient temperature conditioned to low (L) [100-200], medium (M) [100-200], high (H) [100-340] values of the component, as well as the whole-range correlation (WR) [0-350].

Use case 3: study of the non-periodic components

The obtained NMF components do not necessarily show recognizable patterns in the calendar or weekly views, even when they are filtered by their largest values, as in the use case 1. This is the case of the less relevant components, such as the 18-th component shown in Figure 3.27. The calendar heatmap of this component shows no recognizable patterns that can explain the learned peak of consumption at midnight that appears when the operation $\text{Agg}_{\text{avg}(\text{comp}18)} \Pi_{\text{Hour}} C$ is visualized in the barchart at the middle of the figure. These non-periodic components can be explained by focusing the analysis on those days and hours when the component shows high importance.

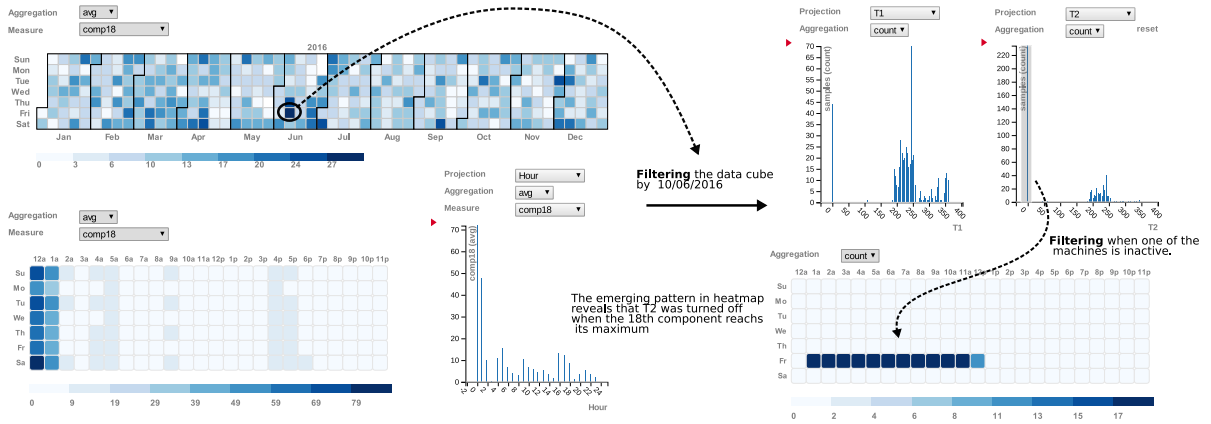


Figure 3.27: The exploration of the 18th component as an example of analysis of a non-periodic component. The user focus the analysis on one of the highlighted days, looking for evidences in the rest of the attributes that can explain the causation behind the component.

Thus, through filtering operations, the view will be specialized to this component and the user can search recognizable patterns in other attributes’ projections. Thus, in the example of Figure 3.27, the calendar view highlights the 10th of June as a day where the 18-th component has a large influence. After filtering the data cube by this day, we observe in the histograms of T1 and T2 (configured by means of the operation $Agg_{count} \Pi_{T1} C$ and $Agg_{count} \Pi_{T2} C$) that during this day both machines were working.

After filtering when T2 is inactive (brushing on zero kW bar), the week heat map reveals that T2 was turned off at midnight just when the peak occurs. When T2 is switched off, the cooling machine T1 had to supply the cooling demand alone, raising its consumptions, and producing the peak of consumption learned by the 18-th component. In a similar way, the 18-th component seems to highlight winter days where a machine of type P is turned off instead of a T machine.

In the light of these results, the peak of the 18th component can be associated with the increment of the total consumption due to the shutdown of a cooling machine. With this case of use, our VA paradigm has not only assessed the non-periodic NMF factors, but also proved the combination of NMF and the data cube allows the user to focus the attention on specific scenarios in a fluid way.

Use case 4: going deeper into the anomalies

During the interactive analysis of the NMF components, some of them show relevance

in unusual days, whose neighboring days do not seem relevant for that component. The principal hypothesis is that in these days a novelty occurs, and the NMF decomposition tries to reproduce it by increasing the importance of the component that fits the best to the anomaly.

In a similar way than in the former use case, we can explore in detail these abnormal days, looking for the cause of the novelty in the views of the other attributes. In Figure 3.28, the third component highlights the 2nd of May as one of these abnormal days. After filtering by this day, the projection of the cube on the P1 reveals that the P1 machine was turned on during this day. By selecting the values of consumption greater than zero, the week heatmap shows that P1 was chilling until 4:00 p.m., and then it was replaced by other machine.

In the projections of the P1 and T1 machines, and after selecting in them the values greater than zero, the week heatmap reveals that two machines were turned on at the same time. The control program which regulates the chilling of the hospital sent the start-up order to T1 and the stop command to P1, but something goes wrong, because P4 was also started and it was chilling during only one hour.

By this abnormal start-up of P4, we have demonstrated that the NMF factors help to “narrow down” the analysis resulting in an effective way to spot abnormal behaviors in the HVAC system of the hospital, that can be explained in detail by specializing the view through data cube filters.

3.5 Conclusions

This chapter presents a novel energy load monitoring approach suitable for large non-commercial buildings, where an unsupervised NILM model, data visualization techniques and interaction mechanisms are successfully integrated into a visual analytics approach, which enables the user to incorporate his prior knowledge to the analysis and to find relevant correlations between hidden patterns and context factors.

The *non-negative matrix factorization* (NMF) algorithm have been used to disaggregate the main energy consumption from a hospital complex, obtaining interpretable individual components that consist in a parts-of-a-whole and sparse representation of the daily electric power demand. Following our VA analytics formulation, the hyperpa-

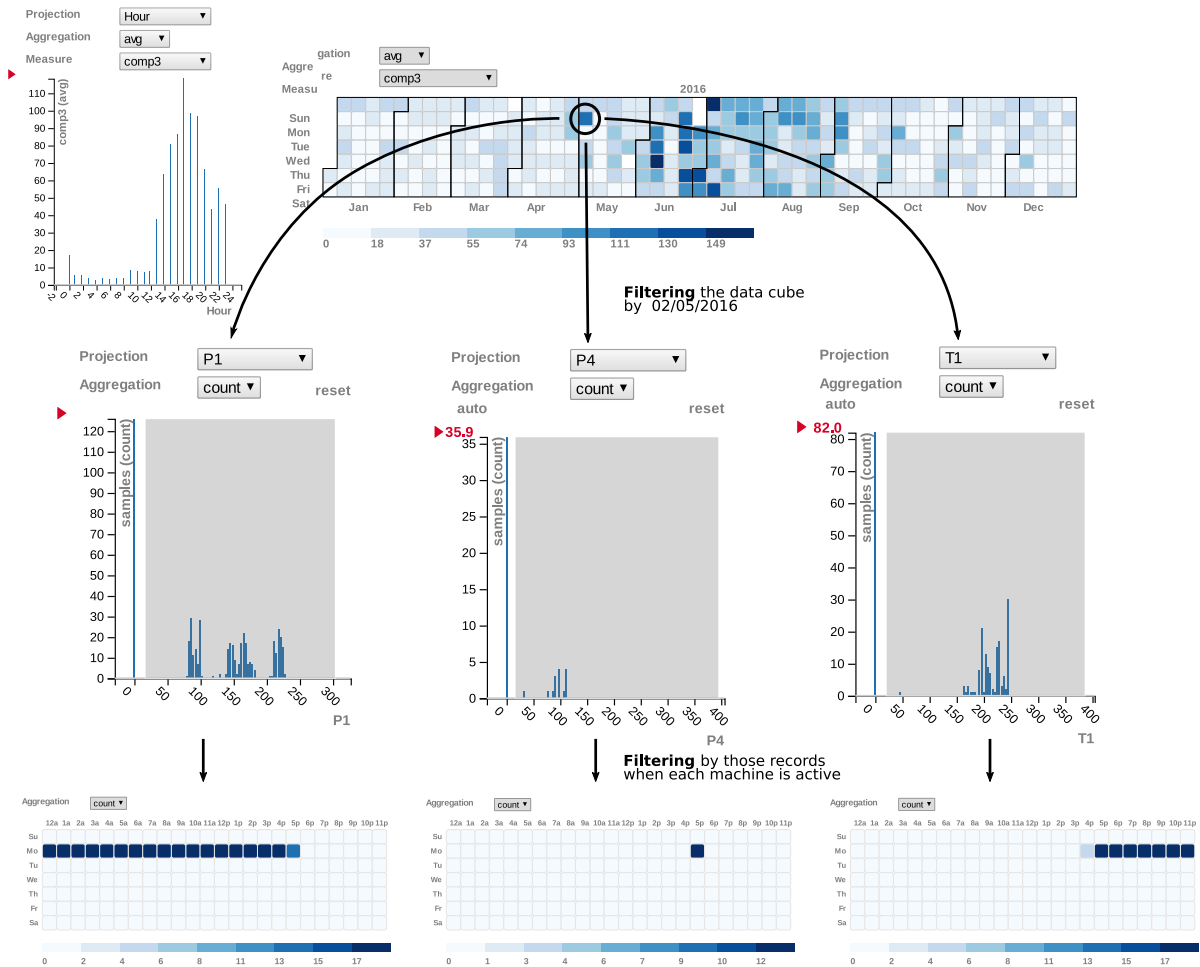


Figure 3.28: Searching for the causes of the anomalies. After selecting the 2nd of May, an anomaly in the starting of the cooling machines is discovered by means of filters in the individual consumptions of the cooling machines.

parameters of the NMF, such as the number of epochs or the number of the components of the decomposition are user-driven parameters in the suggested dashboard.

It is shown that a large number of NMF factors are needed in order to achieve an accurate decomposition, reducing the interpretability of the decomposition. In order to enhance the understanding of the resulting NMF components, the NMF outcome together with several context factors of the facility are arranged in a *data cube structure* that speeds up filter and selection interaction mechanisms in data. The data cube approach enables users to set and modify a broad number of filters on the context factors and the NMF components, obtaining an immediate visual feedback in the used data visualizations techniques, which allows to modulate the analysis from a general overview

of the NMF components to an analysis in detail of an specific anomaly highlighted by a NMF component.

These visualizations have been specifically designed to support the data cube output and its operations, following the visual principle designs. Thus, interactive barcharts, heatmaps and sparklines are suggested to represent the data cube output.

This results in a human-in-loop analysis, where users can visually connect the NMF factors, context factors and their prior knowledge in insightful visualizations, contextualizing the resulting energy disaggregation.

NILM APPROACHES FOR LARGE BUILDINGS BASED ON DEEP LEARNING TECHNIQUES

Despite the suggested NMF-based energy disaggregation model extracts highly interpretable latent components from main consumptions of buildings and enables users to naturally associate them with recognizable events or systems in the proposed VA approach, it does not provide users with a high-resolution feedback, since a complete disaggregation capable of identifying the exact waveforms of each individual consumption is not computed. A rich energy feedback prompts users to adopt energy-saving behaviors [154, 155] so that complete NILM disaggregation techniques, such as the HMM-based or DNN-based models mentioned in chapter 2, are also convenient approaches in energy monitoring tools for large buildings.

Most of the aforementioned state-of-art methods have been designed and tested in residential facilities and only a few of them have tackled the problem of NILM in non-residential large buildings [46, 114, 115]. At the beginning of chapter 3, we have suggested that the differences between non-residential and residential facilities could lead to an inaccurate performance of the state-of-art residential models, when they are applied to non-residential scenarios. Hence, the applicability of cutting-edge NILM models in non-residential facilities is not straightforward and it should be further studied.

The aim of this chapter is to study how to transfer cutting-edge NILM models,

originally developed for residential facilities, to non-residential installations, using as a benchmark the data from the Hospital of Leon described in section 3.1.1. We will identify which aspects of the NILM models, specially those from the DNN-based models, should be adapted to large buildings, and we will suggest an accurate approach suitable for large facilities called *fully-convolutional denoising auto-encoder* (FCN-dAE).

In the literature of NILM for large buildings, the feedback about how energy is being spent is assumed as the main goal, but the properties of the large submetering facilities may suggest other potential uses of the NILM systems, which can also increase the energy efficiency or aid in the daily maintenance tasks. These potential extra functionalities will be explored through real use cases in the hospital, which will rely on the estimated disaggregation.

Following our problem formulation, the NILM operation and the obtained disaggregation will be denoted as D and $\mathbf{p}(t)$, respectively. Although the study of an accurate NILM model and its potential utilities within a large facility is by itself a contribution, the underlying paradigm of our work is to integrate NILM approaches in VA tools which promote a fluid user-machine analysis loop. Bearing this in mind, our NILM approach should satisfy the following requirements to enable the interaction with the user:

- It should provide a set of specifications S_D whereby the user can modulate its outcome according to his mental map, hypotheses or intentions.
- In order to make the user \leftrightarrow NILM interaction intuitive, the computation of a new $\mathbf{p}(t)$ after a change in S_D should be as fluid as possible.

In this chapter, alternative mechanisms capable of modulating the behavior of DNN-based energy disaggregation models and producing meaningful alterations in $\mathbf{p}(t)$ will also be explored with the aim of providing more intuitive and fluid user \leftrightarrow NILM interactions.

In the remainder of this chapter, the related work will be explored in section 4.1. In section 4.2, the NILM models based on *denoising Auto-Encoders* will be particularized to the NILM problem for large buildings, and our FCN-dAE will be presented. A novel mechanism of interaction between the user and the DNN-based NILM models will be suggested in section 4.3. Finally, in section 4.4, all the ideas and models suggested will be tested on the hospital data.

4.1 Related work

Since the early days of NILM, HMM-based models [8, 80, 99] have been among the most studied approaches in the literature to obtain a high resolution disaggregation $\mathbf{p}(t)$. In these models, the main active power of the facility $P(t)$ is considered the observable result of hidden (not directly observed) internal states of the appliances throughout the time. Hence, a discrete HMM model is defined by the following elements:

- A set of $S = \{S_1, S_2, \dots, S_N\}$ possible *internal states*.
- A *transition matrix* $\mathbf{A} = \{a_{ij}, 1 \leq i, j \leq N\}$ that defines the probability of transitioning from an internal state S_i to S_j :

$$a_{ij} = P(q_{t+1} = S_j | q_t = S_i)$$

where q_t represents the active internal state at time t , and $a_{ij} \geq 0$, $\sum_j a_{ij} = 1$

- An *emission matrix* $\mathbf{B} = \{b_i(o_k), 1 \leq i \leq N, 1 \leq k \leq L\}$ representing the probability that a state S_i generates an observation o_k from a known vocabulary \mathbf{V} , so that $o_k \in \mathbf{V}$. In our energy disaggregation problem, \mathbf{V} can be a continuous or discrete dictionary of consumptions, and $b_i(o_k)$ is a probability density function.
- An initial probability distribution over states $\boldsymbol{\pi} = \{\pi_i\}$, where π_i is the probability that the sequence will start with the state S_i :

$$\pi_i = P(q_1 = S_i), 1 \leq i \leq N$$

The parameters of the HMM models $\lambda = \{\mathbf{A}, \mathbf{B}, \boldsymbol{\pi}\}$ are estimated by the *Expectation-Maximization* algorithm, using the observable sequence, and the sequence of internal states is inferred by the *Viterbi* algorithm, making use of dynamic programming.

In the case of NILM, we can assume that a sequence of main consumptions $\mathbf{P} = \{P_1, P_2, \dots, P_T\}$, is the result of the emissions from a *Markov chain* of states $\mathbf{q} = \{q_1, q_2, \dots, q_T\}$, where the state at a time step only depends on the state at the previous time step (Markov property):

$$P(q_{t+1} | q_1, q_2, \dots, q_t) = P(q_{t+1} | q_t)$$

In [8], the authors extended the initial HMM approach to *Factorial Hidden Markov Models* (FHMM) [156], where the sequence of internal states is replaced by a set of independent sequences, one per appliance, as it is shown in Figure 4.1. The new internal states $\mathbf{q} = \{\mathbf{q}^{(1)}, \mathbf{q}^{(2)}, \dots, \mathbf{q}^{(m)}\}$ is a set of sequences $\mathbf{q}^{(m)} = \{q_1^{(m)}, q_2^{(m)}, \dots, q_T^{(m)}\}$ representing the states of the appliance m throughout the time. Thereby, the observed sequence of main consumptions \mathbf{P} is the sum of the emissions from each independent sequence of appliance states.

In [8], it was also introduced three variants of the vanilla FHMM: 1) *Conditional Factorial Hidden Markov Model* (CFHMM), where the probability of transitioning from a state is conditioned by context variables, such as time of the day or day of the week; 2) *Factorial Hidden Semi-Markov Model* (FHSMM) where the probability of transitioning between states also depends on the state occupancy durations; and 3) *Conditional Factorial Hidden Semi-Markov Model* (CFHSMM) where the extensions in 1) and 2) are combined into the same model. The authors reported that the inference of the exact hidden states is intractable by means of dynamic programming algorithms, such as Viterbi approach, so that they suggested estimating the states by *simulated annealing* optimization model.

In [99], the *Additive Factorial Approximate MAP* (AFAMAP) algorithm was proposed as a variant of the simulated annealing for inferring the internal states and the observed sequence of main consumptions, being a robust and convex approach for estimating the observations from the states in a FHMM model. Authors reported average accuracies about 87%, using FHMM models with AFAMAP.

Despite the high performance of HMM-based methods, they also present several drawbacks. In [71], the following drawbacks of NILM HMM-based methods are enumerated.

- *Not all appliances satisfy the Markov property.* The estimation of the next state of some appliances (specially type II) may involve long-term dependencies, rather than just the previous state.
- *HMM-based models do not explicitly model state duration.* The duration of states in a standard HMM model is determined by means of a geometric distribution [8], which implies that the probability of an appliance remaining in the same state exponentially decreases with each time step. This could affect to the performance in




Figure 4.1: FHMM diagram extracted from [8].

appliances with long state duration (such as TV or air conditioning). In [8], authors addressed this problem in their CFHSMM, where prior state duration distribution were imposed in the state transition. Additionally, several authors [157, 158], out of NILM field, suggested mechanisms to integrate the state duration to the HMM models, but all of them increase the computational time or implementation complexity of the resulting models.

- *Some HMM-based models require that all appliances are included in the model.* HMM models are specially sensitive to outliers or energy loads that are not attached to any sequence of the internal states. In real facilities, some devices, which present low energy demand and are sporadically turned on, are excluded from the training datasets, and therefore it is hard to know the exact number of internal state sequences, which may decrease the performance of HMM-based models.
- *The input of HMM-based models is typically preprocessed.* Many works apply noise-removal computations, normalizations or resampling operations over the input main consumptions. In these operations, fine-grained appliance-specific information, that can be useful for the appliance identification, could be dropped.

To address these disadvantages, in [42, 64, 71] the authors suggested applying *deep*

neural networks (DNN) to raw and low-frequency main consumption measurements, motivated by the rise of DNNs in other fields, such as image recognition or text analysis. In these fields, DNN-based models have been able to extract complex and meaningful features from raw data (e.g., images) in a hierarchical way, so that learned features in a layer are a combination from the features of the previous layer, getting more abstract features throughout layers [159].

This automatic extraction of “*hierarchical features*” has prompted the authors in [64, 71] to adapt DNN-based approaches to the NILM problem. Their principal motivation was to investigate if the automatic feature extraction provided by DNNs results in more robust to noise and accurate approach than hand-engineering features, and if the resulting features can be valid for NILM in a wide range of appliances and facilities (transferability). Following our definition of the energy disaggregation operation in section 1.2:

$$F_m : P(t) \rightarrow p_m(t) \quad (4.1)$$

the DNN-based models address the function F by means of a neural network:

$$\mathbf{p}_m = D_{m,\theta}(\mathbf{P}) \quad (4.2)$$

where each appliance-specific load sequence $\mathbf{p}_m = \{p_m(1), p_m(2), \dots, p_m(T)\}$ of size T is extracted from the main consumption sequence $\mathbf{P} = \{P(1), P(2), \dots, P(T)\}$ of the same size by means of the network D_m , whose layer’s weights θ are specifically optimized for the appliance under study. Thereby, a complete NILM system for a facility with m appliances would be a set of trained networks $\{D_{1,\theta}, D_{2,\theta}, \dots, D_{m,\theta}\}$, that is, one *sequence-to-sequence* model per appliance to be disaggregated, as it is graphically shown in Figure 4.2.

In the first steps of DNN-based NILM, *convolutional neural networks* (CNN) [17, 84, 160], *denoising auto-encoders* (dAE) [85] and *recurrent neural networks* (RNN), such as *Long Short-Term Memory* (LSTM) [86], were proposed as architectures for the $D_{m,\theta}$ models. In [71], the authors evaluated these architectures, using data from the five houses of the *UK-DALE repository* [1], and they concluded that the DNN-based models outperform the previous HMM-based approaches, specially the dAE and RNN

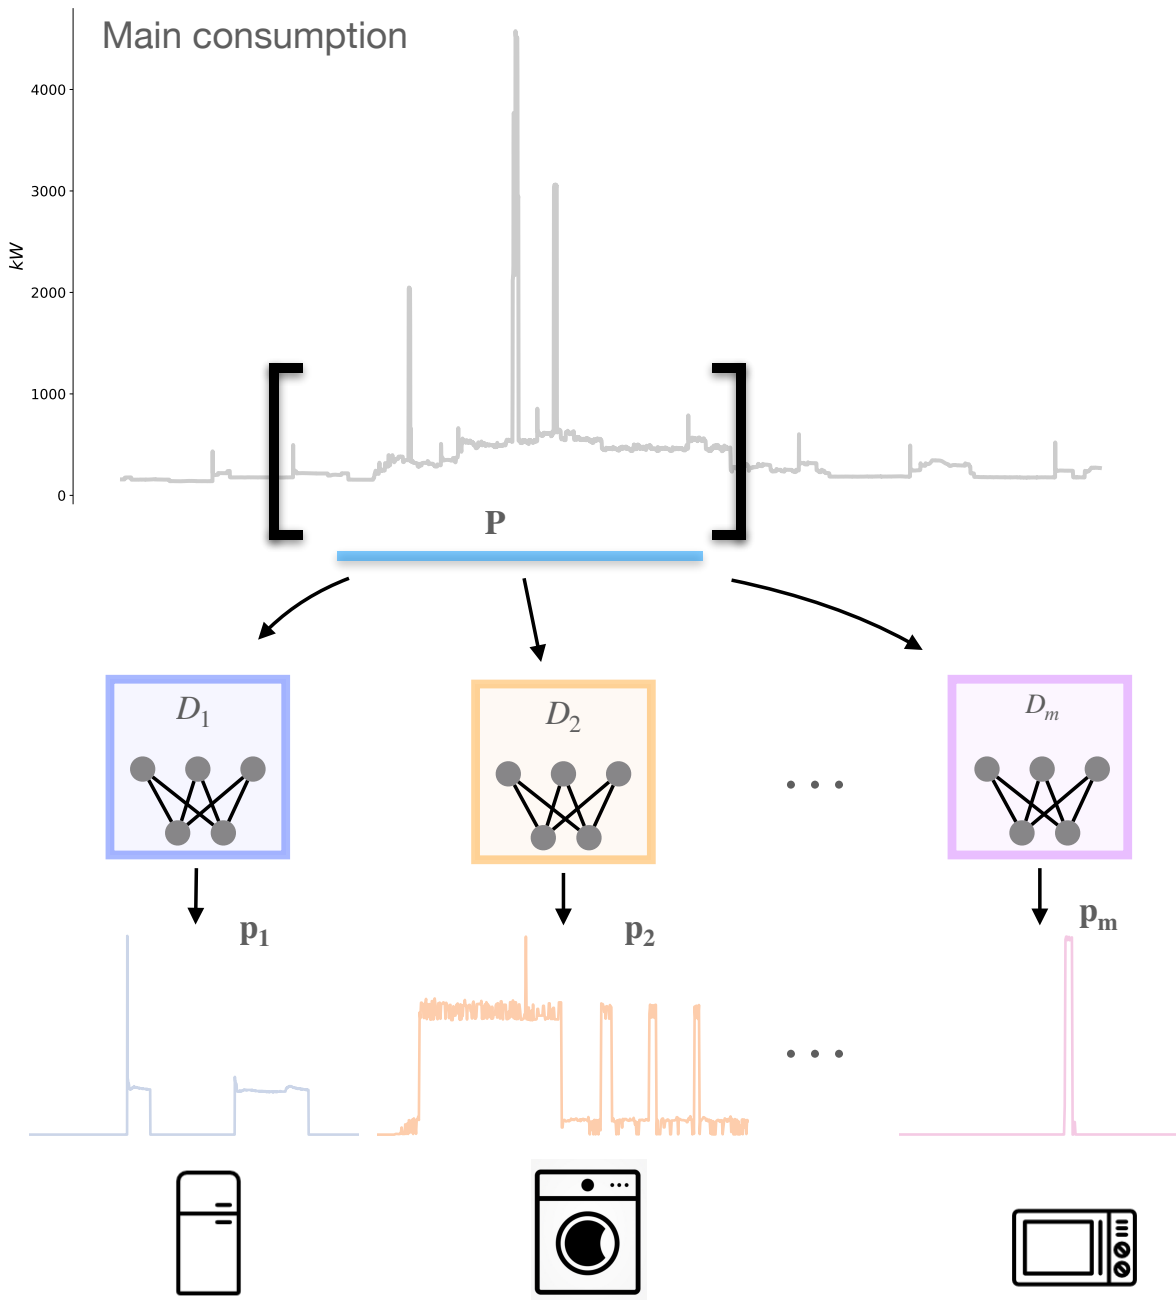


Figure 4.2: Intuition of sequence-to-sequence DNN-based NILM models.

architectures. They also pointed out that the dAE models generalize better for unseen scenarios (houses that are not used in the training stage).

Subsequent research on DNN-based NILM approaches have confirmed that the models based in RNN, such as *bidirectional LSTM* [94] or *Gated Recurrent Units* (GRU) [95], and dAE approaches [88, 96–98] are accurate techniques for NILM. In [161], the

authors have suggested a novel *sequence-to-point* approach, in which a CNN-based $D_{m,\theta}$ estimates a single point $p_m(t)$ of the individual consumption m from a sequence of the main consumption \mathbf{P} , instead of the complete sequence \mathbf{p}_m , showing that the sequence-to-point paradigm can outperform both dAE approach from [71] and AFAMAP models.

In [88], the feasibility of a generic CNN-based model for different types of appliances, using the same hyperparameters, has been explored. In order to improve the transferability, the authors in [111] have suggested a generic bidirectional LSTM capable of dynamically adapting the model to the changes in context factors (e.g., seasonality, weather or different appliance’s operating modes), by monitoring the deviations in accuracy of the model and retraining the entire system, when these deviations rise. All these works demonstrate the enormous attention that DNN-based NILM models have recently attracted, and their potential interest to solve the open issues in NILM.

Despite the aforementioned progresses, the DNN-based methods have been evaluated and designed for residential facilities, and only a few of them [162, 163] have been applied to large non-residential buildings. In [162], the authors proposed a LSTM network to forecast the consumptions of several educational buildings with a resolution of 60 minutes. Beyond the ability of learning features in time series from energy demand of large buildings, we seek a disaggregation with a higher resolution that allows the user to identify *switch-on events* of a few minutes. In [163], an introspection of the activations and weights of a *convolutional Auto-Encoder* trained to reconstruct the main consumption of the Hospital of León was presented. By simple visualizations of the weights and activations of the intermediate convolutional layers, complex patterns of consumption and relations between different areas of the hospital are insightfully discovered. Despite the suggested introspection helps the user to gain knowledge about the electrical installations, it does not provide a complete energy disaggregation.

In this work, we propose to tackle NILM for non-residential buildings by means of DNN-based models, focusing our attention on the dAE architectures, due to their simplicity and their stability in the training phase. RNN-based architectures show similar accuracies in NILM problems, but they are difficult to train because of unstable and vanishing gradient problems [164, 165]. For these reasons, we will start adapting dAE architectures similar to the one proposed in [71]. In addition to transfer dAE-based NILM approaches to large buildings, we will also study useful functionalities of

DNN-based models in other energy management aspects by means of several use cases.

4.2 Denoising Auto-Encoder models as NILM systems in large buildings

Considering the previous discussion about DNN-based NILM methods, in this section, we initially analyze the state-of-art dAE architecture [71] and then, some of its particularities, when it is adapted to consumptions from large buildings, will be discussed. Taking into account these particularities, we will present a novel dAE-based approach, more convenient for large buildings than vanilla dAE, which is based on *fully-convolutional neural networks*. Advantages of the fully-convolutional approaches will be also analyzed in terms of robustness when they are deployed in real installations.

4.2.1 Vanilla denoising Auto-Encoders (dAE)

Denoising Auto-Encoders (dAE) were initially proposed in [85] to extract robust features from input data that are partially corrupted. Authors showed that dAEs extract more structured and discriminating features in high-dimensional data (e.g., images) when they are trained to reconstruct partially destroyed inputs.

Although they were designed to initialize the weights of deeper DNNs in applications, such as image recognition, the dAE-based approaches can be also applied to the NILM field, by reformulating the energy disaggregation problem as a denoising task. Hence, D_m functions can be considered as noise removal functions that map a “noisy” main consumption sequence \mathbf{P} to a “clean” appliance-specific sequence \mathbf{p}_m . Therefore, if we consider the sequence \mathbf{P} as a sum of the sequences \mathbf{p}_m from the downstream individual consumptions

$$\mathbf{P} = \sum_{m=1}^N \mathbf{p}_m \quad (4.3)$$

then the energy demand from each individual load \mathbf{p}_m could be defined as the main consumption minus the remaining demand

$$\mathbf{p}_m = \mathbf{P} - \sum_{l=1, l \neq m}^N \mathbf{p}_l \quad (4.4)$$

Taking this into account, the term $\sum_{l=1, l \neq m}^N \mathbf{p}_l$ is considered the noise to be removed by the dAE function D_m

$$D_m : \mathbf{P} \rightarrow \mathbf{p}_m \quad (4.5)$$

The dAE architecture is based on a standard *Auto-Encoder* (AE) approach, which is trained to reconstruct the input into the output as close as possible. In order to avoid that the AE model degenerates into the identity function, the input is forced to be mapped into an intermediate and compact representation, before generating the output. This intermediate representation is typically called *bottleneck* or *latent space*, and it forces the model to identify the most relevant features of the input and to discard irrelevant information.

In the particular case of the energy demand data, an AE can be trained to reproduce the input sequences $\mathbf{P}^{(i)}$ of main consumptions into the output, firstly mapping them into the latent space by the *encoder* e_θ , and then reconstructing the output estimation $\hat{\mathbf{P}}^{(i)}$ from the latent space by means of the *decoder* d_ϕ . The trainable weights θ and ϕ of the encoder and the decoder, respectively, are optimized by minimizing a similarity loss function. Typically, the *mean squared error* (MSE) is minimized in AEs for all training examples i . The similarity function is also known as the *reconstruction error*, and the minimization process can be formally defined as:

$$\begin{aligned} \theta^*, \phi^* &= \arg \min_{\theta, \phi} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\mathbf{P}^{(i)}, \hat{\mathbf{P}}^{(i)}) \\ &= \arg \min_{\theta, \phi} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\mathbf{P}^{(i)}, d_\phi(e_\theta(\mathbf{P}^{(i)}))) \end{aligned} \quad (4.6)$$

Note that the reconstruction error is represented as \mathcal{L} , and it is formulated as $\mathcal{L}(\mathbf{P}^{(i)}, \hat{\mathbf{P}}^{(i)}) = \|\mathbf{P}^{(i)} - \hat{\mathbf{P}}^{(i)}\|^2$, if the MSE between the input and its reconstruction is used as similarity function.

In the case of a dAE, the aforementioned AE-based architecture is trained to remove an undesired noise or information from the input $\mathbf{P}^{(i)}$, getting in the output the individual consumption sequence $\mathbf{p}_m^{(i)}$. A key difference of dAEs with respect to AEs is that the input and output sequences differ, so that the optimization of θ and ϕ is now:

$$\begin{aligned}\theta^*, \phi^* &= \operatorname{argmin}_{\theta, \phi} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\mathbf{p}_m^{(i)}, \hat{\mathbf{p}}_m^{(i)}) \\ &= \operatorname{argmin}_{\theta, \phi} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\mathbf{p}_m^{(i)}, d_\phi(e_\theta(\mathbf{P}^{(i)})))\end{aligned}\tag{4.7}$$

where the reconstruction error is computed between the output estimated sequence and the target input sequence without the undesired information, which leads the model to reconstruct a different sequence from the input sequence. We can now model D_m in (4.5) as a dAE, which learns to estimate the individual consumptions:

$$\hat{\mathbf{p}}_m = D_m(\mathbf{P}) = d_{\phi, m}(e_{\theta, m}(\mathbf{P}))\tag{4.8}$$

Once the principles of the dAE approach and its application to the NILM problem are explained, a specific topology has to be defined. We suggest a first approach based in [71] as a base line in our subsequent experiments. In Figure 4.3a, this topology is detailed, where the encoder projects the input sequence \mathbf{P} into the latent space by means of a 1D convolutional layer followed by two fully-connected layers. The decoder reconstructs the estimated $\hat{\mathbf{p}}_m$ from the latent space by means of two fully-connected layers followed by a 1D convolutional output layer with only one filter. All the hyperparameters that define the layers of the encoder and decoder (i.e., units, kernel size and number of kernels) shown in Figure 4.3a were empirically set to those values which produce the best performance in terms of reconstruction error.

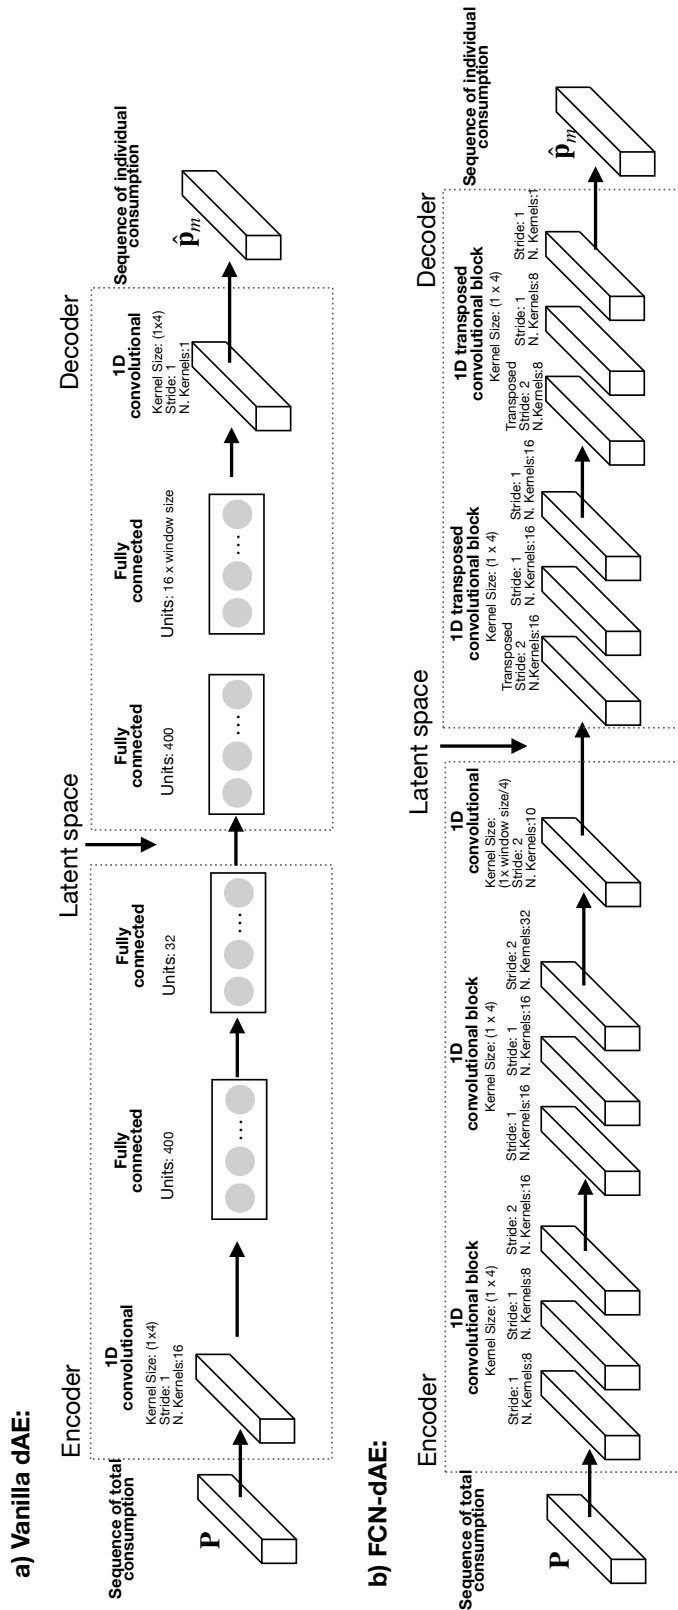


Figure 4.3: The studied dAE-based topologies. a) Vanilla dAE architecture where each cell symbolizes a layer explained by the number of units in the case of fully-connected layers; or the kernel size, stride and number of kernels in the case of convolutional layers. b) Proposed FCN-dAE architecture made entirely by 1D convolutional layers and 1D transposed convolutional layer.

4.2.2 Disaggregation process and the input window size selection

The aforementioned dAE models D_m are typically trained with pairs of sequences $(\mathbf{P}, \mathbf{p}_m)$ both with duration L , so that the trainable weights are optimized by minimizing the loss function \mathcal{L} using modern optimization algorithms, such as the *Adam optimizer* [166], that provide efficient implementations of the gradient descent algorithm.

Once models D_m are correctly trained, our set of denoising functions can infer sequences of individual consumptions. The main consumption sequence to be analyzed is often significantly longer (e.g. one-week or one-month sequences) than the training sequences of duration L , so that it must be split into segments of duration L suitable to be processed by D_m .

In [71, 97], authors suggested splitting the sequence by sliding a window of size L along the main consumption under study, with a stride of M between contiguous windows. This process is illustrated in Figure 4.4, where the parameters L and M define the input temporal context used by D_m and determine how centered is a sample $P(t)$ within the resulting segment. The position of $P(t)$ in the input segment has a direct impact on its corresponding individual consumption estimation $\hat{p}_m(t)$, since it determines the temporal context before and after in the input.

In order to provide an appropriate context for each estimation $\hat{p}_m(t)$, previous approaches proposed windowing operations with $L > M$, introducing an overlapping between windows so that $\frac{L}{M}$ estimations $\hat{p}_m^{(i)}(t)$ are computed, because the same measurement $P(t)$ is included in several input segments and in different positions. All estimations $\hat{p}_m^{(i)}(t)$ are then aggregated by their average [97] or by other aggregation functions, such as the median [71], to come up with an output sequence of the same length as the input one. Using the average of the estimations, the authors reported accurate estimations of the individual loads, when the overlapping of the windowing is long enough.

The parameter L defines the amount of information that the models use to produce the output sequence. The temporal input context strongly influences the result obtained by the dAE models. Models trained with too long windows are prone to fail in reproducing short-term signatures of the individual consumptions. By contrast, models trained with too short input windows are prone to poorly estimate the individual loads, since the resulting dAE models lack of input temporal information, and long-term dependencies in

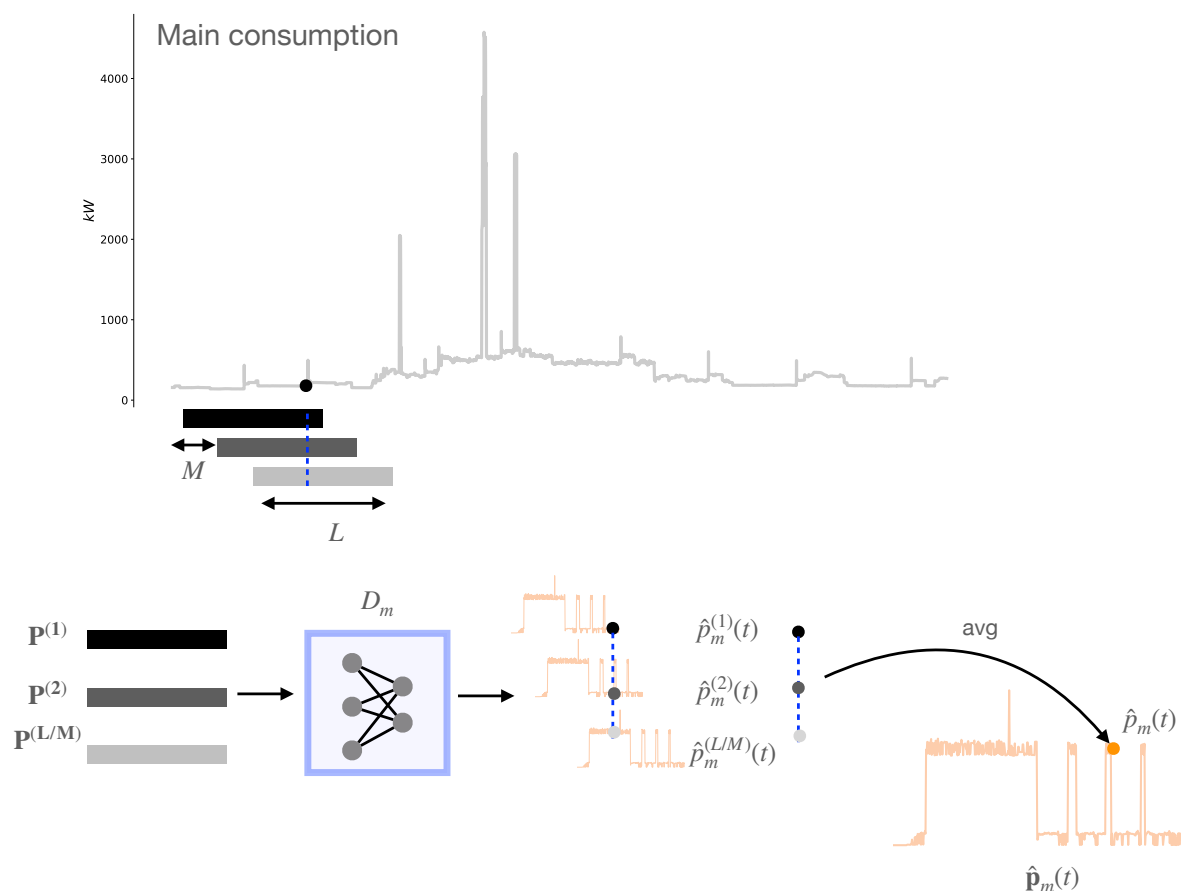


Figure 4.4: Disaggregation process for DNN-based approaches.

the main consumption, that are relevant to the decomposition, are not taken into account during their training. Figure 4.5 exemplifies this effect with a dAE trained with different window sizes. Note that the input sequence length L also determines the first and output layers of the architectures shown in Figure 4.3, and therefore their number of trainable parameters. Hence, long input windows lead to models with a large amount of weights, increasing the risk of overfitting and making the models more specific to the training data, which may reduce their transferability (generalization) to “unseen” scenarios.

In residential NILM, L is usually set appliance-by-appliance so that the average activation of the appliance completely fits in the input window. This rule is easy to satisfy in residential facilities, where the time between on and off events in appliances is usually known and it does not vary, but transferring this rule to large non-residential buildings is not straightforward, since most of the individual loads show a continuous energy demand, as it was shown in section 3.1.2. Hence, an architecture robust to the




Figure 4.5: Different estimations of fridge consumption produced by a dAE model trained for different input sequence lengths.

input window size capable of learning both long-term and short-term dependencies, and of disaggregating all the individual loads with the same configuration is highly desirable in NILM for non-residential buildings.

4.2.3 Fully-convolutional denoising auto-encoder (FCN-dAE)

Taking into account the nature of consumptions in large buildings and the issue of the parameterization of dAE-based models, we enumerate the following requirements that a DNN-based model applied to non-residential building has to satisfy:

1. *Robustness*. A model robust to the window size.
2. *Multi-resolution*. A model capable of extracting long-term (low-frequency patterns) and short-term (high-frequency patterns) behaviors from the input main consumption sequences.
3. *Generalization*. A generic model, that with the same configuration (number of layers, number of kernels, kernel sizes, etc.), achieves accurate disaggregations for a broad spectrum of individual consumptions with different natures and behaviors.
4. *Efficiency*. An efficient model in terms of training time, inference time and number of weights that can be easily deployed without requiring a large amount of memory and be integrated in VA applications without losing fluidity.

Most of the aforementioned requirements are shared to the time series modeling approaches. In the literature, the time series analysis is often considered as a *sequence modeling* problem where techniques, such as *Recurrent Neural Networks* (RNN) specially *Long-Short Term Memory* (LSTM) architectures [86] have been widely used [167–169].

Although RNNs are theoretically independent to the input window size and able to learn short-term and long-term features, these architectures also entail several disadvantages, as it is pointed out in [165], where authors empirically showed that the promise of “infinite memory” is compromised because of their unstable gradients (exploding/vanishing gradients) and large memory requirements for training. In this work, the authors also suggested an alternative for time sequence modeling based on *fully-convolutional neural networks* (FCN) [170–172]. Several experiments in sequence modeling tasks suggest that simple FCN models outperform standard LSTM networks, showing better “*memory retention*” and being much more simple architectures.

In the light of these results, FCN-based architectures are good candidates for NILM in large buildings, where the estimation of an individual consumption may depend on long-range features, such as daily patterns. Its simplicity in number of weights and its stable training process may also contribute to improve the transferability and deployment in new scenarios.

FCN architectures were initially defined in [170] as image segmentation models capable of dealing with inputs of arbitrary size. Authors in [170] suggested that, due to the equivariance to the translation of the convolutional layers, FCN models can be considered as “*deep nonlinear filters*” instead of general nonlinear functions resulting from a conventional CNN. In the case of 1D time series, denoting the activation of a layer as \mathbf{y} , and the activation of previous layer as \mathbf{x} , if both activations are connected by means of a convolutional layer with only one kernel and a linear activation function, \mathbf{y} is computed by:

$$\mathbf{y} = c_{m,s}(\mathbf{x}, \mathbf{k}) \quad (4.9)$$

where we refer to $c_{m,s}$ as the convolution operation, and to \mathbf{k} as the kernel convolved over the input with size m and s as the stride between convolutions. Hence, the output at the i -th position

$$y_i = \sum_{i,m} x_{i \times s + m} k_m \quad (4.10)$$

is only path-connected to a few elements $\{x_{i \times s + \delta_i}\}$, $0 \leq \delta_i < m$ of the input \mathbf{x} called *receptive field*. An intuition about this idea is graphically shown in Figure 4.6.

Hence, if a feature in \mathbf{x} is shifted in time, its representation in \mathbf{y} will be also shifted. The concept of that an element of a layer is locally connected with several elements of the input is maintained even when multiple convolutional layers are stacked. A composition of two convolutional layers, results in

$$c_{m,s}^{(1)} \circ c_{m',s'}^{(2)} = \left(c^{(1)} \circ c^{(2)} \right)_{m'+(m-1)s',s's} \quad (4.11)$$

where the an activation y_i depends on $\{x_{i \times s s' + \delta_i}\}$, $0 \leq \delta_i < m' + (m - 1)s'$ elements of the input.




Figure 4.6: Graphical representation of the convolutional layers, receptive field and the composition of convolutional layers.

By contrast, in a fully-connected layer an output unit y_i is connected to the whole input vector \mathbf{x} , aggregating all the input information and losing information about the position and the short-range features. This vanishing of temporal information is further exacerbated when the input size is enlarged, making the results smoother as it is shown in Figure 4.5. FCN models may help to prevent the vanishing of temporal information through layers thanks to the network-wide convolutional nature described in (4.11), preserving both short-term and long-term features.

Bearing this in mind, transforming the aforementioned standard dAE model into a *fully-convolutional denoising Auto-Encoder* (FCN-dAE) could lead to more robust NILM models in terms of input size variations and concurrent long-term and short-term dependencies.

In Figure 4.3b the suggested FCN-dAE is shown, where all the fully-connected (dense) layers of previous dAE model were replaced by convolutional and transposed convolutional blocks. In the case of the encoder, dense layers were replaced by convolutional blocks, where each of them consists of two 1D convolutional layers with stride 1 followed by a 1D convolutional layer with stride 2. Convolutional layers with stride 2 shrink the dimension of the input sequence along the convolutional blocks, making the bottleneck a condensed representation, which is enlarged back by the decoder. The decoder expands the bottleneck representation by applying 1D *transposed convolutional blocks* [173, 174], which consist of a 1D transposed convolutional layer followed by two standard 1D convolutional layers with stride 1.

The configuration of the two 1D convolutional layers before reducing the dimension of the input sequence by means of a convolution with stride 2, and after increasing the size of the bottleneck representation by a transposed convolution was inspired by the well-known VGG image recognition model [175]. This layout makes FCN-dAE model complex enough to tackle the NILM task in non-residential large buildings.

The disaggregation is computed following the same process as in the vanilla dAE models explained in section 4.2.2 and, in the same way as dAE models, the FCN-dAE approach is a sequence-to-sequence transformation that maps a main consumption sequence \mathbf{P} to an individual-load-specific sequence \mathbf{p}_m , being necessary a trained model D_m for each individual appliance. The number and size of kernels were also set to those values that reduce reconstruction error the most.

4.3 Integration of FCN-dAE into visual analytics

Returning to the problem formulation of section 1.2, the NILM operation D should provide the users with mechanisms that modulate its outcome according to the user's hypotheses. This modulation is associated with a set of actionable parameters S_D capable of configuring D according to the user's intentions. In the case of the aforementioned DNN-based methods, the available parameters can be enumerated as follows:

- The M and L parameters from the windowing operation described in section 4.2.2.
- The optimization parameters involved in the training of the D_m models. The type of optimizer, the learning rate or the number of epochs are examples of these parameters.
- The parameters related to the architecture of D_m , such as the number, type and settings of the layers.

However, in order to make a modulation driven by the aforementioned parameters effective, a retraining of the entire DNN model is needed. The training time of D_m models explained above range from minutes to hours (depending on the input sequence length, see section 4.4.4). These training times do not support the paradigm of fluid interaction, and therefore they would slow down the human-machine loop, decreasing the confidence on the results and the sensation of immersion.


For this reason, mechanisms capable of modulating the DNN-based models D_m in a fluid way should be further studied in order to integrate DNN-based NILM models into the proposed VA paradigm.

4.3.1 Conditioning in deep neural networks models

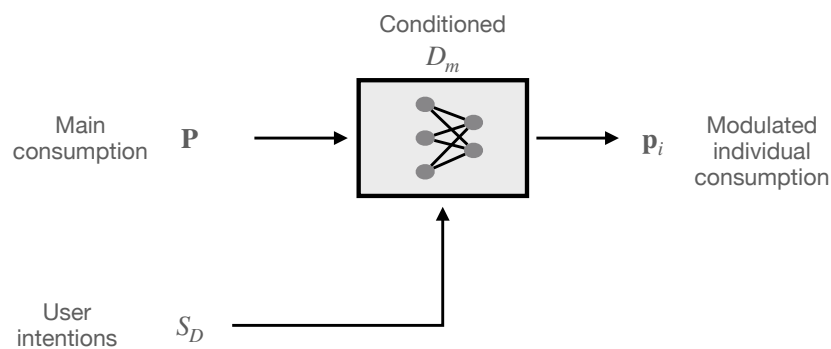
The interaction pathway between the user and the disaggregation operation D can be seen as a fusion between two sources of information from two different contexts:

1. *Energy demand context*, where a sequence of aggregated energy consumption is decomposed into a set of individual consumption sequences.
2. *User context*, where the intentions of the user are indicated.

One way of integrating both contexts is *conditioning* the main context by means of a proper representation of the second context. In this paradigm, the main task is called the *conditioned task*, and the information by which is modulated, is called the *conditioning task*. Both pipelines are graphically represented in the diagram of Figure 4.7a. This idea can be translated to our VA paradigm applied to NILM, as it is shown in Figure 4.7b, where the main task would be a trained DNN-based model, and the conditioning information would be the user intentions, which are represented as the parameters S_D .



(a) General conditioning diagram.



(b) Conditioning diagram within our energy disaggregation paradigm.

Figure 4.7: Graphical diagrams of the idea of conditioning DNN models

The conditioned DNN models have been recently studied, and they have been focused on modulating the intermediate activations of the conditioned network by means of the conditioning information. In [176], the authors have suggested a conditional generative DNN approach based on *Generative Adversarial Networks* [177], where the conditioning information is the class label of the image to be generated. In this model, the conditioning information is directly introduced into the network, concatenating it at the input of the *discriminator* and the *generator* networks.

Concatenating information to the input of a fully-connected layer is similar to bias its output activations, since the matrix product can be divided into a sum of two subproducts (one for the normal input and the other one for the conditioning information) [178]. This kind of conditioning is known as *conditional biasing*, and it is also implemented in the well-known architectures *conditional PixelCNN* [179] and *WaveNet* [171].

Other models use mechanisms of *conditional scaling*, which multiply the intermediate activations by a scaling vector. In this category, several works have suggested the same input as the conditioning information used to modulate subsequent intermediate activations. This kind of models are known as *self-conditioned* models, and they have been used in image recognition models [180] and in natural language processing approaches [181].

Recent *style transfer* approaches [182, 183] enable the user to change the artistic style of a paint, giving to the model an example of the target style (conditioning information). The main contribution of style transfer approaches is a new conditioning method called *conditional instance normalization*, which modulates the output feature maps of several intermediate convolutional layers as follows:

$$\hat{\mathbf{F}}_{j,k}^{(i)} = \gamma_{j,k}^{(i)} \left(\frac{\mathbf{F}_{j,k}^{(i)} - \mu}{\sigma} \right) + \beta_{j,k}^{(i)} \quad (4.12)$$

where the feature map of the j -th convolutional layer and the k -th channel $\mathbf{F}_{j,k}^{(i)}$ computed from i -th input sample $\mathbf{x}^{(i)}$ is firstly normalized by the mean μ and standard deviation σ computed across the spatial axes of $\mathbf{F}_{j,k}^{(i)}$. Then, the normalized $\mathbf{F}_{j,k}^{(i)}$ is scaled and shifted by means of an affine transformation defined by the parameters $\gamma_{j,k}$ and $\beta_{j,k}$. In style transfer models, $\gamma_{j,k}$ and $\beta_{j,k}$ are computed from the conditioning style (conditioning information) using a secondary network, which learns a representation of the target style and generates a proper set of parameters for the affine transformations of all conditioned layers.

The conditional instance normalization approach encouraged the authors in [178, 184] to define the general-purpose *Feature-wise Linear Modulation* (FiLM) approach, shown in Figure 4.8. FiLM conditioning mechanism is inspired in the aforementioned conditional instance normalization, but FiLM gets rid of the normalization from the affine transformation:

$$\text{FiLM}(\mathbf{F}_{j,k}^{(i)}) = \gamma_{j,k}^{(i)} \mathbf{F}_{j,k}^{(i)} + \beta_{j,k}^{(i)} \quad (4.13)$$

Thus, the FiLM modulation can be considered a post-normalization affine transformation, where the affine parameters are computed from the conditioning information a_i by

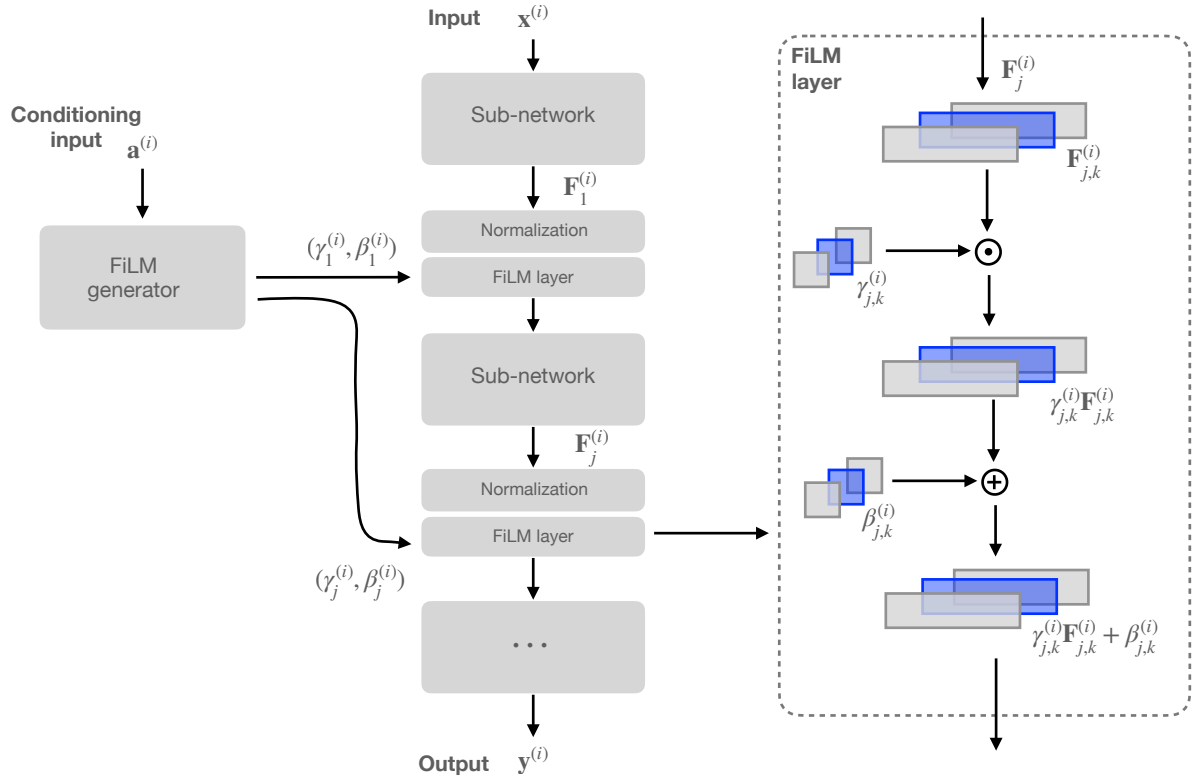


Figure 4.8: FiLM conditioning approach.

functions f_c and h_c :

$$\boldsymbol{\gamma}_j^{(i)} = f_c(\mathbf{a}^{(i)}) \quad \boldsymbol{\beta}_j^{(i)} = h_c(\mathbf{a}^{(i)}) \quad (4.14)$$

Functions f_c and h_c can be implemented by an auxiliary network, and both functions can share parameters, so that the same network produces all pairs of affine parameters $(\boldsymbol{\gamma}_j^{(i)}, \boldsymbol{\beta}_j^{(i)})$ for all the feature maps to be modulated. This secondary network is called *FiLM generator*, and the main neural network on which the FiLM modulation are applied is called *FiLM-ed network*. Examples of the FiLM conditioning, FiLM generator and FiLM-ed network are illustrated in Figure 4.8.

The FiLM framework has been applied to integrate difficult contexts, such as questions and images in visual reasoning tasks [184], demonstrating that by applying simple affine transformations across feature maps, meaningful modulations can be achieved. Additionally, FiLM modulations are computationally efficient both in terms of number of

weights, and in terms of computational time, since the FiLM generator only learns two parameters per channel of each modulated convolutional layer, so that the number of affine parameters is independent from the input size, and FiLM modulation does not entail large delays, as a complete retraining of the network.

For these reasons, the FiLM conditioning approach perfectly fits in our VA paradigm, allowing powerful and meaningful modulations of DNN-based NILM models with a low computational cost.

4.3.2 Multi-task FCN-dAE

Following the FiLM conditioning paradigm, we can modulate our suggested architecture FCN-dAE by defining a conditioning source of information $\mathbf{a}^{(i)}$. In our VA formulation, $\mathbf{a}^{(i)}$ was defined as the user-actionable parameters S_D . A first approach for our conditioned NILM model could be to replace the aforementioned list of sequence-to-sequence models, where a DNN is trained per individual consumption, by a unique sequence-to-sequence DNN that disaggregates the target individual consumption indicated in the conditioning input.

The resulting conditioned model is a *multi-task NILM*, where the user indicates the target individual node by means of S_D . This paradigm assumes that the FiLM-ed network, in this case our FCN-dAE, has learned a rich enough reservoir of features for all the individual consumptions and the FiLM conditioning mechanism enhances the representative features of the target consumption.

Hence, such a multi-task NILM approach entails a conditioning input S_D that indicates the target individual consumption, and it should be interpretable for the user. A set of attributes that defines the individual consumption (e.g., location, type of consumption or any prior knowledge about the target consumption) or a simple *one-hot encoding* that indicates the individual node to be disaggregated are examples of S_D for our multi-task NILM paradigm. In a first approximation of our multi-task NILM paradigm, the target individual encoding is introduced by a one-hot encoding.

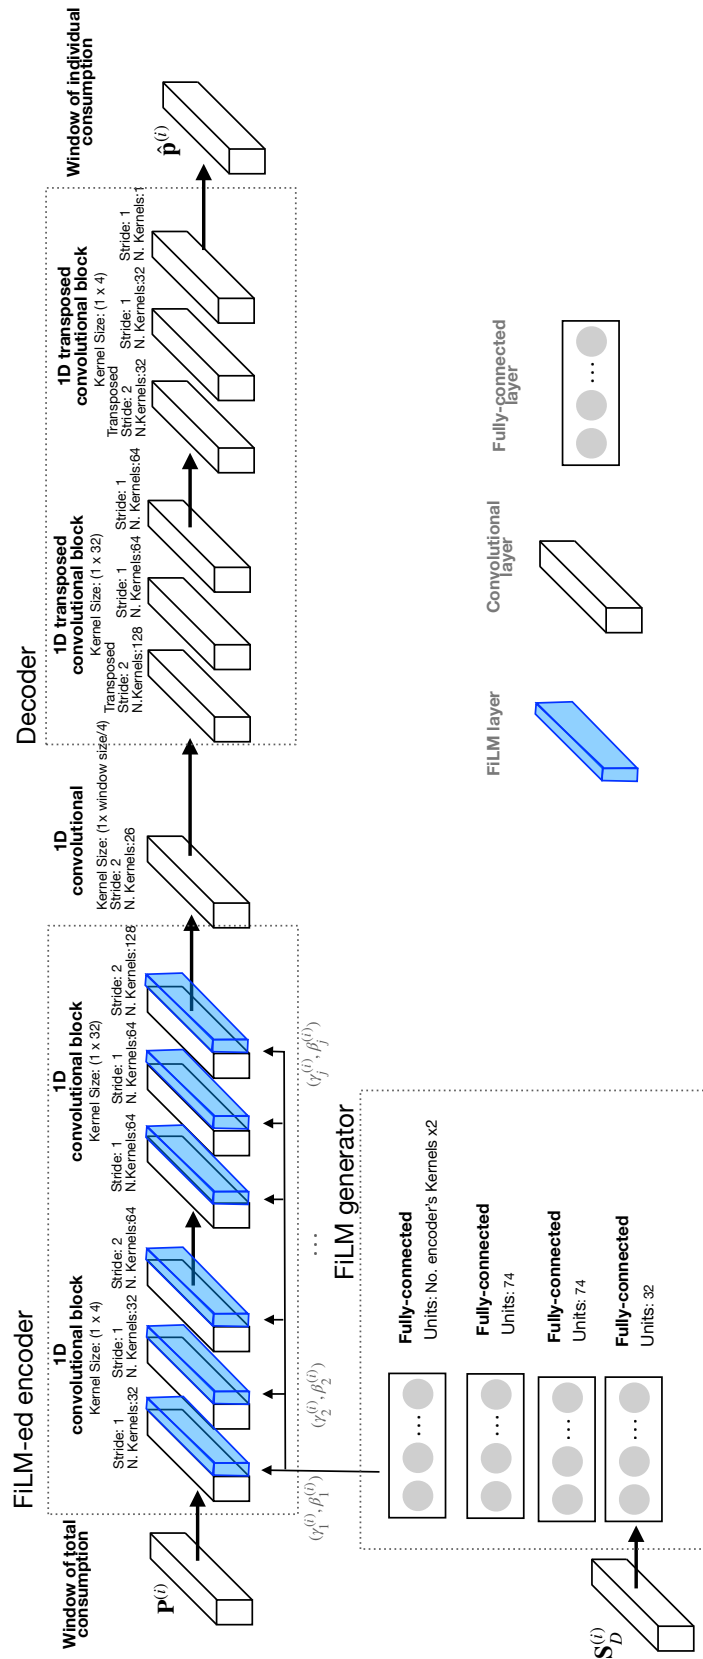


Figure 4.9: Multi-task FCN-dAE architecture.

In Figure 4.9, the suggested multi-task architecture is shown. The conditioning one-hot encoding $S_D^{(i)}$ is processed by a FiLM generator made of three fully-connected layers with ReLU activation function and the output affine parameters $(\gamma_j^{(i)}, \beta_j^{(i)})$ are computed by a set of output fully-connected layers with linear activation; one per feature map to be modulated.

As it is shown in Figure 4.9, all the encoder’s convolutional layers are modulated by FiLM layers. Thus, all the features learned by the encoder are adaptively modulated according to the target consumption. Modulating only the encoder, the latent space representation is conditioned to the target consumption which tends to result in a set of subspaces, one per target consumption. Then, the decoder is capable of reconstructing each subspace to the corresponding individual consumption.

Regarding the FiLM-ed network, we use the same topology as in the previous FCN-dAE approach, but the number of kernels of the convolutional layers has been increased in order to ensure a rich enough reservoir of features to cope with all the appliance with same network.

4.4 Results

In this section, the aforementioned DNN-based NILM methods will be tested using the data of the Hospital of León and the UK-DALE datasets, both described in section 3.1.1. After explaining our experimental setup in section 4.4.1, in sections 4.4.2, 4.4.3, 4.4.4 and 4.4.5 our FCN-dAE approach is tested in terms of robustness, efficiency and potential uses in large facilities. Finally, in sections 4.4.7 and 4.4.6, the performance and the new user↔NILM interaction path of our multitask FCN-dAE approach are evaluated.

Most of the results and conclusions related to the FCN-dAE in the following sections were included in the article [9], which has been published in the journal *IEEE Transactions on Smart Grid*.

4.4.1 Experimental set-up

Although DNN-based models are able to resolve extremely complex tasks, such as image recognition or natural language processing, they are often difficult to train, in part because of the large number of user-defined hyperparameters. Among the hyperparameters involved in the training process are those related to the topology of the network,

such as the number and type of layers, the layer-specific parameters (e.g., number of units, number of kernels or kernel's size), the optimization algorithm or the training data. In subsequent experiments, DNN-based methods under study will be based on the topologies shown in previous sections (see Figures 4.9 and 4.3), where the number and type of layers and all the layer-specific parameters were described in detail.

Regarding the training data, all models were trained with pairs of sequences $(\mathbf{P}^{(i)}, \mathbf{p}^{(i)})$ of main and individual consumptions extracted from the facility. The pairs of sequences are the result of a windowing operation with a window size L equal to the input size of the model to be trained.

Once the pairs of sequences are taken, both the main consumptions sequences $\mathbf{P}^{(i)}$ (input sequences) and the individual consumption sequences $\mathbf{p}^{(i)}$ are normalized by subtracting the mean and dividing by the standard deviation, computed from the whole signal data.

In subsequent sections, the models used in the experiments were trained with pairs of sequences extracted from the UK-DALE and the Hospital's datasets, both described in section 3.1.1. In the case of the Hospital, the main consumptions were collected from all the available data from the node *CGBT-2.Red-Grupo*, except one month which was used as a test data. The individual consumptions were gathered from all downstream nodes from *CGBT-2.Red-Grupo*. This division results in a 90% of training data and a 10% of testing data.

On the other hand, the training and testing sequences from UK-DALE were extracted from *House1*, taking one year for training and two months for testing. In addition, data from UK-DALE are downsampled to 60s resolution in order to equal the sample period in both datasets.

As we mentioned in previous sections, models were optimized by means of the *Adam optimizer* [166] using the reconstruction error between the estimated sequences $\hat{\mathbf{p}}^{(i)}$ and their corresponding ground truth $\mathbf{p}^{(i)}$. The number of epochs during training was set to 100 and in each epoch a batch with 100 pairs of sequences was used to compute the Adam's optimization step.

In the inference stage, the disaggregation is accomplished by splitting the whole sequence of main consumption to be decomposed into segments of size L and a stride

between windows of 5 minutes. Both the final disaggregation and the training computations were executed in a PC with a Linux-based operation system and with a *NVIDIA GeForce RTX 2070 6GB* GPU, which provides us with enough computational power and drastically speeds up the optimization of the models and the subsequent disaggregations.

4.4.2 Input sequence length effects

The robustness of the FCN-dAE approach against variations in the input sequence length was assessed by training several complete sets of FCN-dAE models $\{D_1, D_2, \dots, D_m\}$ using the data of the Hospital of León, with different input sequence lengths: 32, 120, 1440 minutes.

In the comparison, the FHMM model from [105], the RNN model biLSTM [71] and the vanilla dAE were also included in order to compare the effects of the input sequence length in the state-of-art models based in AE, RNN and HMM approaches. In the case of the dAE and biLSTM approaches, the same amount of sets of models with the same input sequence configuration and training data as in the case of FCN-dAE were trained. In FHMM models, only one factorial model is trained, since the input time series is processed as a stream of data and all individual nodes are estimated at once. The number of possible hidden states in FHMM is set to two.

The maximum input sequence length in DNN-based models was set to 1440 minutes (1 day), since larger sequences would cause long training latencies and increase the amount of trainable weights, specially in the case of vanilla dAE, causing problems of overfitting. Bearing this in mind, we have kept the topologies as simple as possible limiting the input temporal context of the models to daily sequences which is also consistent with the intrinsic nature of data from large buildings since, as it was shown in 3.1.2, they show a strong daily seasonality.

Hence, all DNN-based approaches were trained and tested in extracting each individual node from the main consumption of the hospital, as it is shown in Table 4.1. This entails 9 different networks per approach and input sequence length, resulting in a total of 81 DNN-based models, which together with the 9 FHMM model make a total of 90 trained models. The disaggregation performances shown in Table 4.1 are measured by NEP, RMSE and MAE metrics (see section 2.4), after applying each D_m model to the test data from the hospital.

CHAPTER 4. NILM APPROACHES FOR LARGE BUILDINGS BASED ON DEEP LEARNING TECHNIQUES

Table 4.1: Disaggregation performance of NILM models with different input sequence lengths in all individual nodes from Hospital. In the first column a FHMM approach is included in order to compare the DNN-based results with a HMM-based approach. The NEP, RMSE and MAE metrics are used as evaluation metrics.

Architecture Sequence Length	FHMM [105] -	dAE [71]			biLSTM [71]			FCN-dAE			
		32	120	1440	32	120	1440	32	120	1440	
MAE	CGBT-2.Montante0	29.9647	2.7193	2.6321	2.4863	2.8637	2.6456	2.9095	2.8171	2.6995	2.5609
	CPD	5.4205	3.4693	3.6887	4.3056	3.4550	3.7128	4.2315	3.436	3.6267	3.9156
	Plantas_2-7	5.5465	2.7913	2.4544	2.4512	2.7887	2.6219	2.4502	2.7519	2.4916	2.109
	Plantas_8-13	3.4196	3.3945	3.0415	2.5463	3.3266	2.8858	2.3686	3.3363	2.8763	2.0707
	Radiologia1	18.1471	6.8708	7.4706	7.0405	6.9443	7.2621	8.2969	6.7407	6.4656	5.9346
	Radiologia2	2.9016	3.0980	3.0451	3.1002	2.9440	2.9278	3.3974	2.9886	2.7157	2.7061
	RehabilitacionA	8.0446	1.2547	1.2237	1.2164	1.2742	1.2592	1.4721	1.2475	1.2945	1.2789
	RehabilitacionB	30.2364	3.6580	3.3045	3.0455	3.7315	3.4512	3.8967	3.6388	3.3004	2.8464
	Subcentral3	5.0117	2.7963	2.6028	2.5553	2.7558	2.6183	2.7142	2.7945	2.6738	2.552
	NEP	CGBT-2.Montante0	0.8261	0.0750	0.0726	0.0685	0.0790	0.0729	0.0802	0.0777	0.0744
CPD		0.0994	0.0636	0.0676	0.0789	0.0633	0.0681	0.0776	0.063	0.0665	0.0718
Plantas_2-7		0.1601	0.0805	0.0708	0.0707	0.0805	0.0757	0.0707	0.0794	0.0719	0.0609
Plantas_8-13		0.1482	0.1471	0.1318	0.1103	0.1441	0.1250	0.1026	0.1445	0.1246	0.0897
Radiologia1		0.4991	0.1890	0.2055	0.1936	0.1910	0.1997	0.2282	0.1854	0.1778	0.1632
Radiologia2		0.2347	0.2506	0.2464	0.2508	0.2382	0.2369	0.2749	0.2418	0.2197	0.2189
RehabilitacionA		0.3832	0.0598	0.0583	0.0579	0.0607	0.0600	0.0701	0.0594	0.0617	0.0609
RehabilitacionB		0.7592	0.0918	0.083	0.0765	0.0937	0.0867	0.0978	0.0914	0.0829	0.0715
Subcentral3		0.1859	0.1037	0.0965	0.0948	0.1022	0.0971	0.1007	0.1037	0.0992	0.0947
RMSE		CGBT-2.Montante0	31.1776	3.6313	3.5056	3.5161	3.8179	3.5716	5.2149	3.7655	3.6545
	CPD	5.6987	3.8705	4.1057	4.6968	3.8690	4.1550	4.6070	3.8444	4.0798	4.3564
	Plantas_2-7	6.7983	3.6113	3.1049	3.0926	3.5647	3.3727	3.1213	3.5407	3.1687	2.7183
	Plantas_8-13	4.1875	4.0551	3.6574	3.1191	4.0464	3.5803	3.0489	4.0385	3.5712	2.6963
	Radiologia1	20.776	10.0057	11.0173	9.8098	9.8136	10.5808	11.6352	9.6737	9.2198	8.771
	Radiologia2	5.1926	5.6233	5.5897	5.6539	5.4948	5.4982	6.0392	5.5192	5.3487	5.3551
	RehabilitacionA	8.8492	1.6526	1.6236	1.648	1.6731	1.6712	2.1273	1.653	1.7090	1.6881
	RehabilitacionB	32.4982	4.7261	4.1982	4.0777	4.8306	4.4338	6.4429	4.7132	4.2299	3.7579
	Subcentral3	5.8894	3.5432	3.315	3.2491	3.4968	3.3588	3.5735	3.5499	3.4090	3.2287

Observing the results, FCN-dAE models yield the best results across metrics and for most of the individual nodes. Furthermore, in most of cases the performance of FCN-dAE improves when the length of the input sequence is increased, being the FCN-dAE models with longest input sequences the best models in all cases except for three individual nodes. Regarding biLSTM approaches, they seem to be specially sensitive to the input sequence length, since it is difficult to set a generic input length for all the nodes and there is not an evident correlation between an increment in the input sequence length and the accuracy. The FHMM approach cannot deal with continuous consumptions with only two hidden states and only performs well in nodes such as *CPD* or *Radiologia2* whose energy remains nearly constant throughout the day.

In general terms, dAE-based models perform similar to FCN-dAE, but in nodes where short-term and long-term features are involved at the same time, their accuracies drop

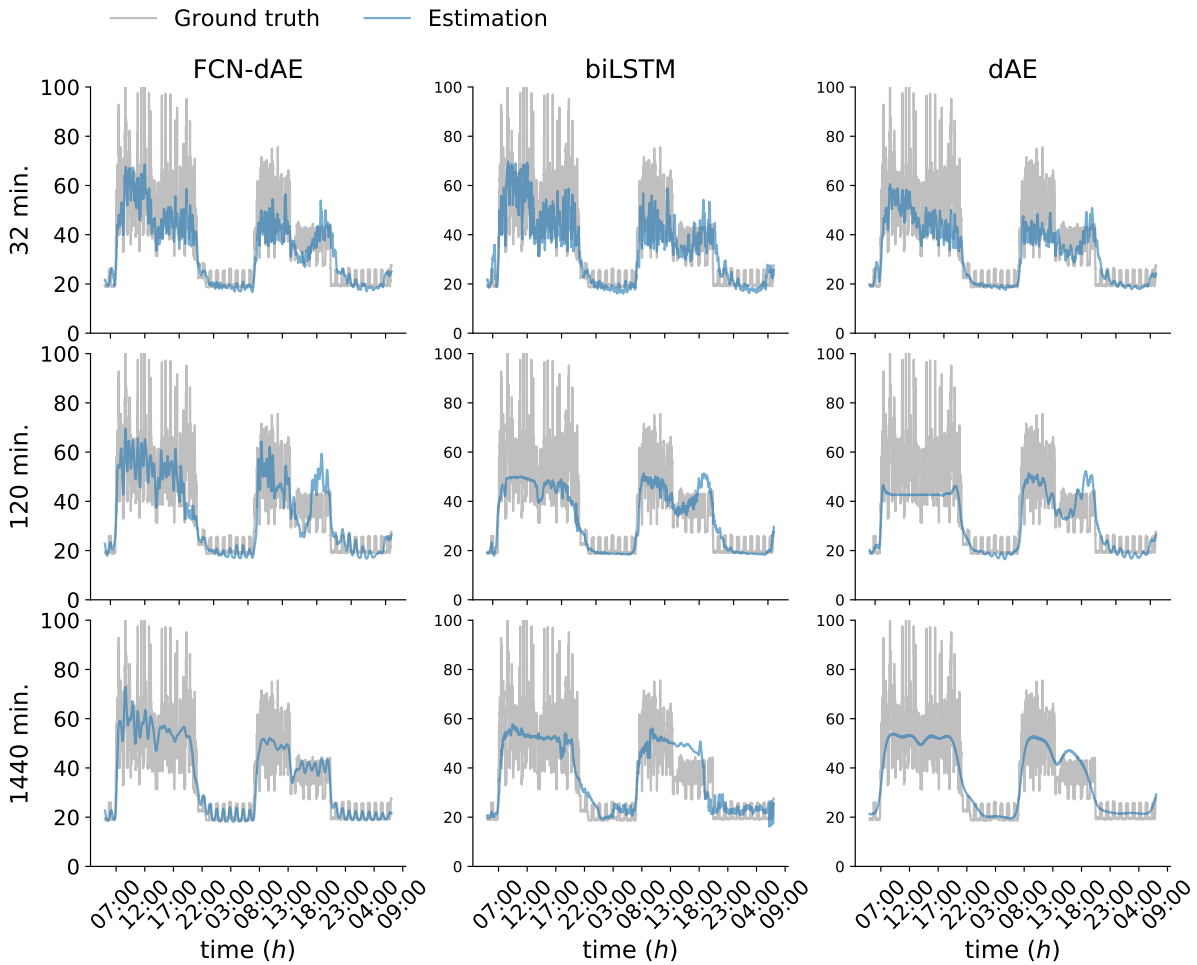


Figure 4.10: Estimations for the node *Radiologia1* obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.

when the input window size is augmented. This is exemplified in node *Radiologia1* of Figure 4.10. In this node, a long-term daily hump together with a short-term consumption during night, are involved. In this scenario, the vanishing of short-term signatures in the dAE models' outcome, when the input sequence length is increased, can be observed.

By contrast, FCN-dAE models are capable of reproducing long and short-term signatures even with an input sequence length of 1440 minutes. This is consistent with the decrease in performance of dAE models with long input sequences observed in Table 4.1, which may be caused by the inaccurate estimation of the short-term behaviors. A similar example is shown in Figure 4.11, where the FCN-dAE disaggregation for node *RehabilitacionB* is more detailed than the estimations provided by dAE and biLSTM.

The fine-grained estimations produced by FCN-dAE bear rich information that could improve the correlation of the disaggregation with events in devices of interest.

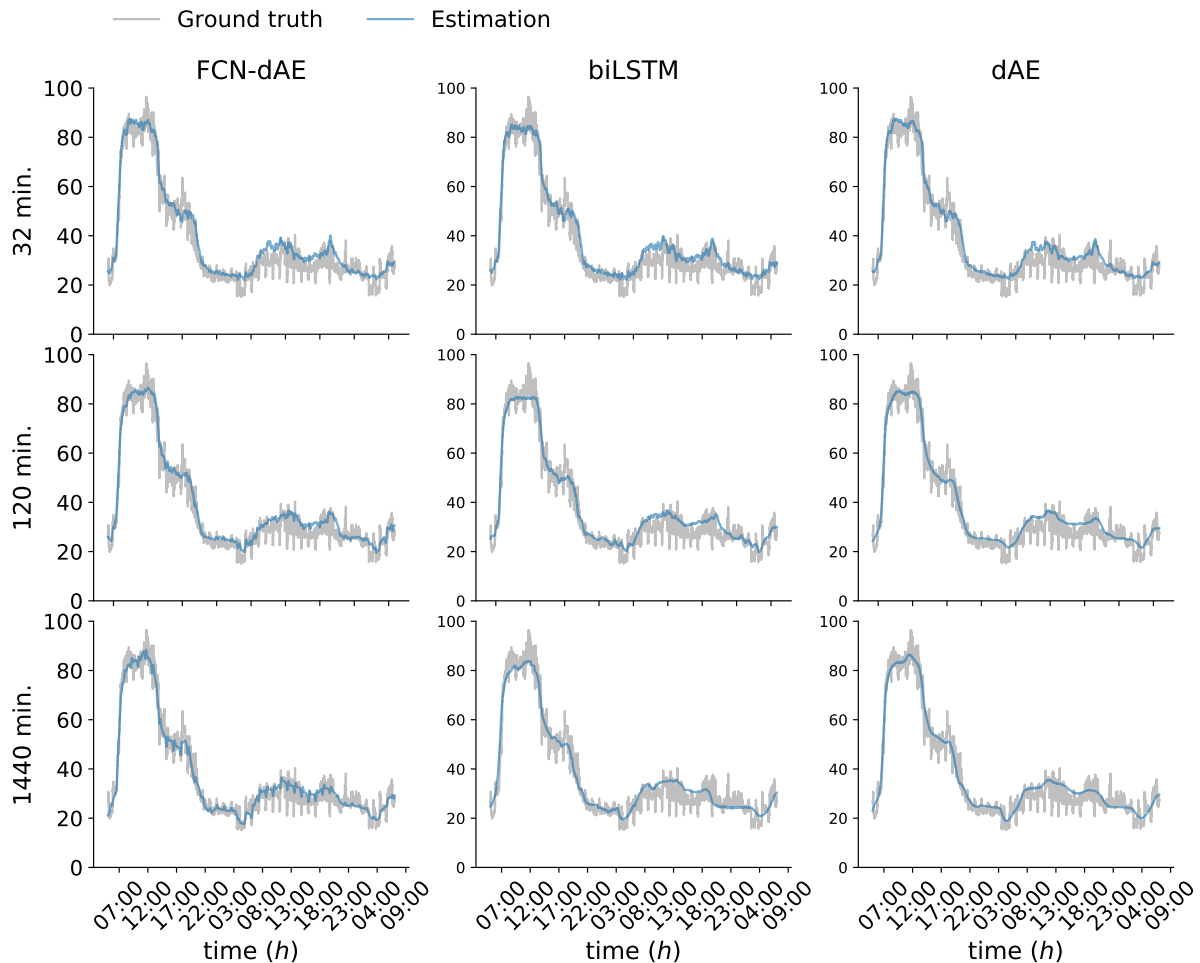


Figure 4.11: Estimations for the node *RehabilitacionB* obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.

A further study about the ability to reproduce at the same time short-term and long-term signatures is shown in Figure 4.12, where the performance of the resulting disaggregations of *Radiologia1* are evaluated by analyzing the residuals between the real and the estimated individual consumptions in terms of frequency. Once the residuals are obtained, their *power spectral density* (PSD) is computed by means of the *Welch method*, that estimates the residuals density in the frequency spectrum, highlighting those frequency bands which concentrate more error. This frequency representation provides us with an insightful evaluation about both short-term and long-term errors.




Figure 4.12: Power spectral density of the residuals calculated from the disaggregations resulting from applying the dAE and FCN-dAE architectures, trained with input window sizes of 30, 120 and 1440 minutes, to the individual node *Radiologia1*. The highlighted harmonics correspond to the errors made by the models at daily (freq. = 0.041 hour^{-1}) and at 75 minutes (freq. = 0.8 hour^{-1}) periods.

In the PSD analysis of Figure 4.12, the residuals obtained from the node *Radiologia1* show two principal harmonics at daily and 75-minutes periods, respectively. Those periods are strongly correlated to the natural behavior of node *Radiologia1*. As it can be seen in Figure 4.10, the *Radiologia1* energy demand has a strong daily seasonality and a periodic load during the night with a duration of about 75 minutes. The error associated with the daily harmonic decreases for all models when the input sequence is enlarged, but the error of the 75-minutes harmonic grows for the dAE models with long input sequences. However, it remains low for FCN-dAE models except for the case of 32 minutes input sequences, where the error may stay high due to the lack of temporal input context.

In the light of these results, the FCN-dAE approaches, when the input length is long

enough, are not only able to outperform dAE models, but also more robust to the input sequence size by accurately reproducing both long-term and short-term signatures in individual nodes.

Finally, the nodes where FCN-dAE approaches perform poorly are those which do not have a broad spectrum of frequencies and can be modeled with smooth estimations. *RehabilitacionA* is an example of these nodes whose energy demand is a flat consumption without large variations of energy, as it is shown in Figure 4.13

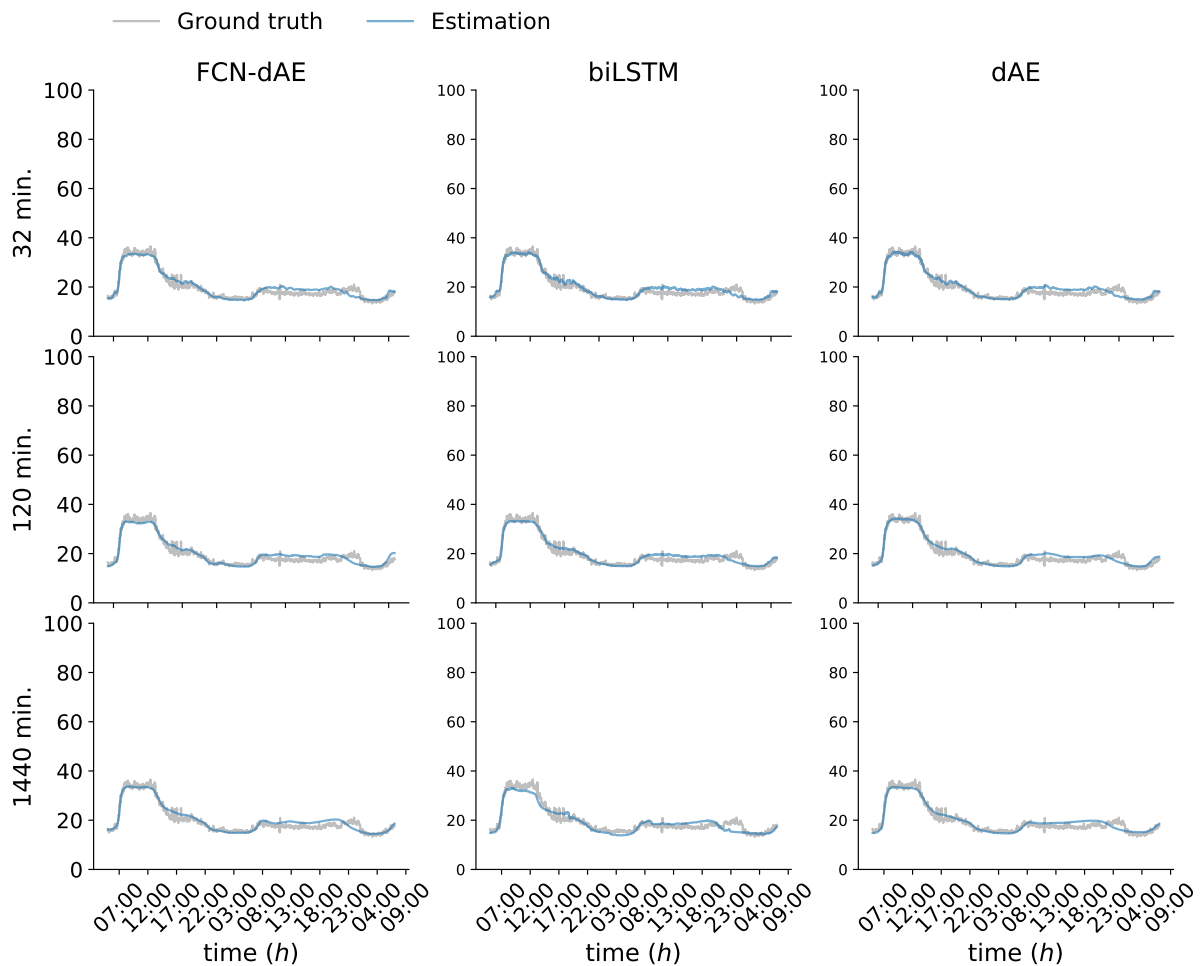


Figure 4.13: Estimations for the node *RehabilitacionA* obtained after applying all DNN-based architectures, trained with input window sizes of 30, 120 and 1440 minutes, to a daily snapshot of the total consumption.

Table 4.2: Disaggregation performance of models with different input sequence lengths in all individual nodes from Hospital. In the first column a FHMM approach is included in order to compare the DNN-based results with a HMM-based approach. The NEP, RMSE and MAE metrics are used as evaluation metrics.

		FCN-dAE				dAE [71]				Seq2Seq [161]				Seq2Point [161]				bLSTM [71]			
		30	60	90	120	30	60	90	120	30	60	90	120	30	60	90	120	30	60	90	120
MAE	FR	13.378	10.365	9.315	8.514	16.603	16.045	16.731	20.431	14.214	11.013	10.743	10.076	13.203	9.994	9.545	10.721	11.091	8.454	8.327	8.551
	WR	7.026	4.748	3.997	3.739	12.438	11.320	10.522	11.281	7.832	6.240	5.763	5.433	7.113	6.552	6.352	6.768	5.536	6.318	13.835	10.209
RMSE	FR	23.951	20.804	20.410	20.123	26.985	26.211	27.376	30.322	24.179	21.063	20.874	20.743	25.415	21.686	20.593	20.222	27.831	23.122	22.456	23.011
	WR	51.334	41.489	37.213	34.216	70.383	65.784	65.778	66.789	53.992	48.391	47.836	43.325	67.465	60.057	54.888	53.340	62.246	62.034	86.833	84.791
NEP	FR	0.276	0.213	0.192	0.175	0.342	0.330	0.345	0.421	0.293	0.227	0.221	0.208	0.272	0.206	0.197	0.221	0.228	0.174	0.172	0.176
	WR	0.342	0.231	0.195	0.182	0.605	0.551	0.512	0.549	0.381	0.304	0.281	0.264	0.346	0.319	0.309	0.329	0.269	0.308	0.673	0.497

4.4.3 Performance of FCN-dAE in residential facilities

The reproducibility and reliability of our approach were also assessed applying the FCN-dAE architecture on a different scenario with individual consumptions distinct from hospital’s ones. Hence, we replied the previous experiments on consumptions from UK-DALE repository, whose house’s individual consumptions are sparse, strongly attached to occupant patterns of consumption and their seasonality is less predominant.

The trained models were the same DNN-based approaches as in the case of hospital experiment, together with the *seq2seq* and *seq2point* approaches recently suggested in [161], which are a CNN-based sequence-to-sequence model and a CNN-based sequence-to-point approaches, respectively.

As in the case of the hospital, for each approach a set of models, one per appliance and input sequence length, were trained. The *fridge* and the *washing machine* were selected as the devices to be disaggregated, since they are two of the most representative appliances in a house and concentrate the majority of the activations and the energy demand. Fridge (FR) and washing machine (WM) are appliances of type I (on/off devices) and II (multi-state devices), which differ from the hospital’s individual consumptions considered of type IV (continuously active systems). By considering these appliances into the evaluation, we are testing the FCN-dAE model for a broad spectrum of different consumptions.

Regarding the input sequence length, we trained all approaches with input sequences of 30, 60, 90 and 120 minutes. Note that the maximum amount of temporal context is drastically reduced, compared to the maximum of 1440 minutes in the hospital, because the ON durations and features are not as long in time as in the hospital, where the performance of the resulting disaggregation may depend on features quite far in time.

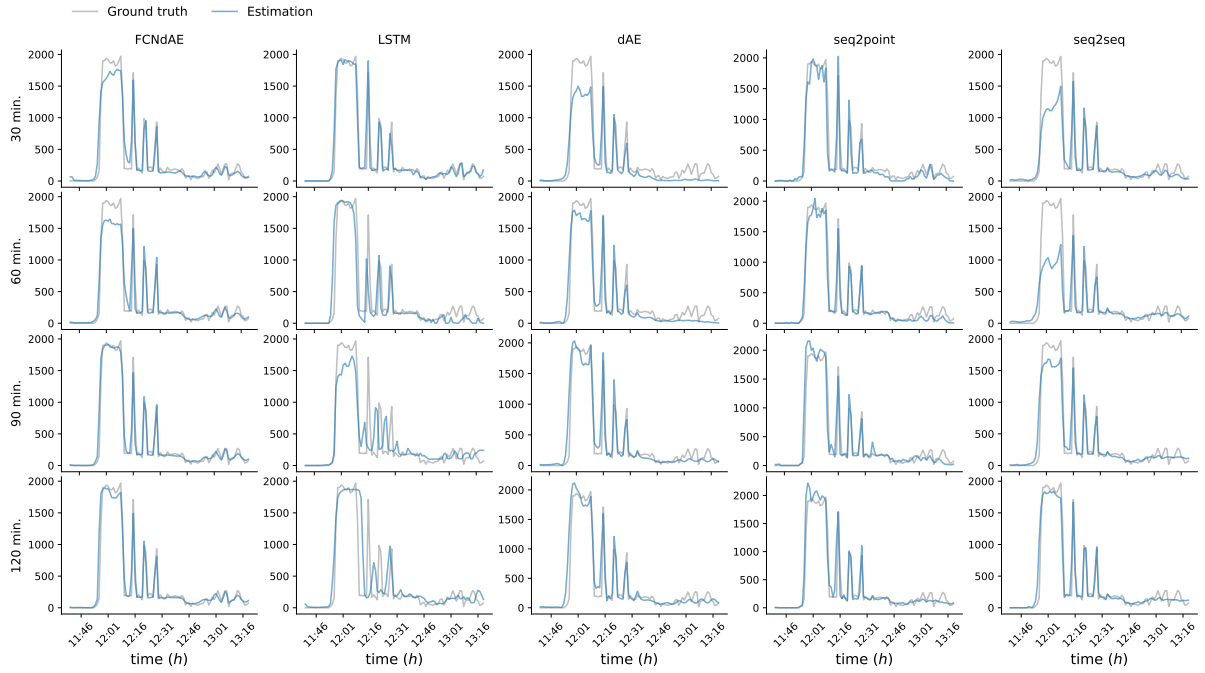


Figure 4.14: Estimations for the appliance *washing machine* (WM) obtained after applying all DNN-based architectures, trained with input window sizes of 30, 60, 90 and 120 minutes.

In table 4.2, MAE, RMSE and NEP metrics of the trained models for the proposed input lengths are shown. FCN-dAE seems to be the most stable approach against input length variations, outperforming all models in the case of the WM, and competing with the biLSTM models with input length of 90 minutes in the case of the FR. In biLSTM models, an increase of the input length does not imply an improvement in the accuracy of the disaggregation. In fact, for input lengths of 120, biLSTM models decrease their performance especially in the case of the WM, as it can be also observed in Figure 4.14.

Regarding dAE models, they seem to be more sensitive to input variations in appliances of type I and II than in continuously active individual consumptions, specially in the case of type I devices (fridge), where the performance sharply drops when the input length is increased. This can be also seen in the corresponding disaggregation, shown in Figure 4.15.

4.4.4 Computational efficiency of FCN-dAE

Another aspect to consider in DNN-based NILM models, specially if they will be deployed in a real scenario, is the computational efficiency. The computational efficiency in DL is

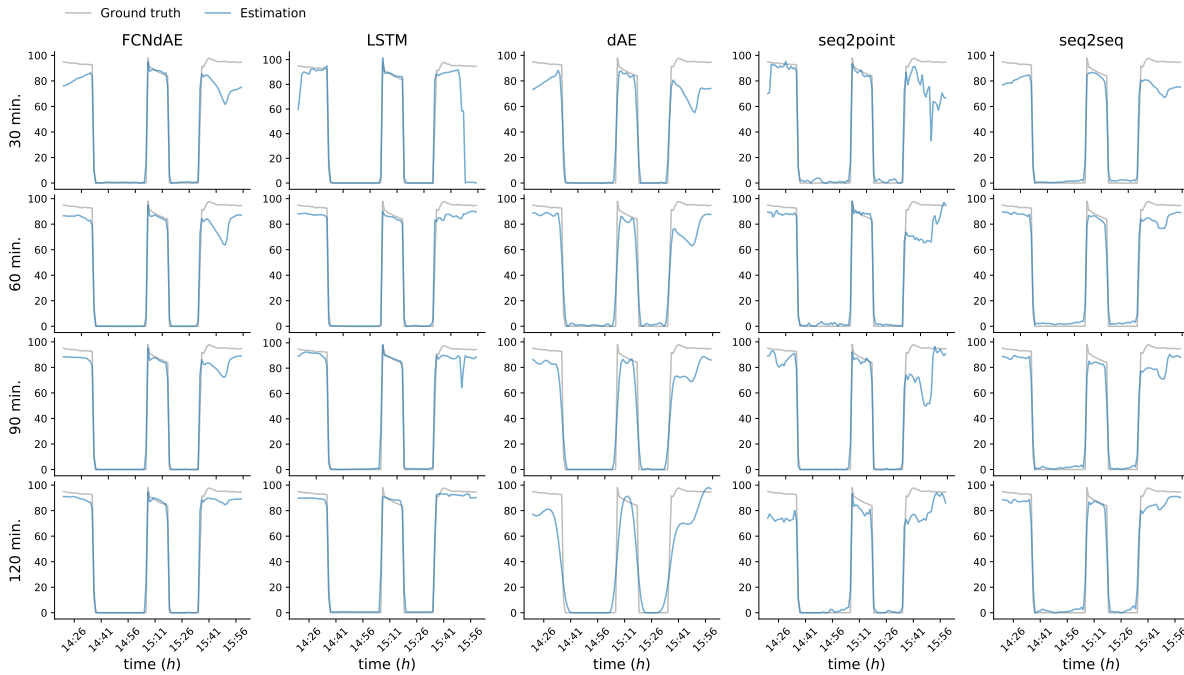


Figure 4.15: Estimations for the appliance *fridge* (FR) obtained after applying all DNN-based architectures, trained with input window sizes of 30, 60, 90 and 120 minutes.

commonly assessed through two factors: 1) the time needed to complete the disaggregation of all individual nodes, and 2) the number of trainable weights. The number of trainable weights defines the minimum needed memory of the device where the NILM system is to be deployed and, on the other hand, the disaggregation speed is a critical feature when outcome of the NILM system is integrated in a VA-based energy monitoring system, where a fluid interaction is a requirement.

Hence, the models under study are analyzed in terms of efficiency, summing the total number of trainable weights and measuring their computation time when they are applied to a single input sequence with the same length as their first layers. Times were computed and measured in a standard laptop in order to assess the efficiency of the approaches in common devices, without expensive GPUs.

Figure 4.16 shows that, although the number of layers is higher in the FCN models, their amount of trainable weights is significantly less than in dAE and biLSTM models. The reason why the FCN models need less parameters is the *parameter sharing* principle of the convolutional layers [159], which makes them more weight-efficient than the fully-connected layers.




Figure 4.16: Computational efficiency in terms of time of computation and number of weights per topology. The first plot depicts the total amount of trainable weights of all the suggested models. Second plot indicates the time in seconds required for each model to compute the estimation from one window of total consumption. Note that the axes in both graphs are logarithmic.

The second row of Figure 4.16 also shows that biLSTM approaches take longer for them to compute the disaggregation of an individual consumption, making them less suitable for systems where a real time prediction is required. The differences in time between dAE and FCN models are minimal in all input lengths, and both architectures need less than 16ms to produce a new estimation, that is the limit for a fluid interaction [126].

4.4.5 Use cases in large buildings

In this section, other potential uses of DNN-based NILM systems for energy management tasks in large buildings will be discussed. Unlike residential facilities, the feedback provided by NILM systems in large buildings may be exploited by experimented maintenance personnel that are in charge of the correct operation of the electric installation. Among their duties are the detection of errors and the maintenance of the submetering systems, in which the NILM system can contribute with useful information that may help them in their daily tasks. In the remainder of the section, we will study how the NILM outcome can be processed to help in tasks, such as the maintenance of the submetering

data acquisition system, filling gaps in the collected data or detecting novelties in the individual circuits which could be caused by potential faults in the facility.

Replacement of real meters

Beyond any doubt, the most evident application of the NILM outcome is the replacement of real measurements from meters by their corresponding estimations $\hat{\mathbf{p}}_m$ from the model D_m that can act as a “virtual sensor”. Note that the temporary replacement of meters can be extremely useful during meter revisions, since the lost measurements during the check-up or reconfiguration processes can be replaced by the NILM outcome, keeping a continuous system monitoring.

In complex facilities with a large number of circuits and subsystems, a complete monitoring of all the systems would involve an unaffordable submetering set-up. For this reason, some non-critical circuits of the facility are typically excluded from the submetering architecture. Thanks to the replacement of real meters in other circuits, the maintenance personnel could install the unplugged meters in these non-measured areas in order to gain knowledge about them without any additional cost.

Gap filling

A well-known issue in real facilities is the occasional loss of data caused by malfunctions of the meters, the network traffic or issues in data storage systems, such as databases. The aforementioned NILM models D_m are able to accurately replace lost measurements in individual consumptions by their corresponding estimations $\hat{\mathbf{p}}_m$. Hence, the measurements of \mathbf{p}_m corresponding to the gap can be directly filled with their estimations $\hat{\mathbf{p}}_m$.

A graphic example of gap filling is shown in Figure 4.17, where a gap in node *RehabilitacionB* is filled with the corresponding estimations provided by dAE and FCN-dAE models, both with input sequences lengths of 1440 minutes. This procedure can be applied as long as the main consumption \mathbf{P} remains uncorrupted, and the model D_m is available.

Nevertheless, if the failure also has an impact on the main consumption, causing data loss in the sequence \mathbf{P} , the input segments of D_m will be affected, and therefore the accuracy of the estimations $\hat{\mathbf{p}}_m$ will be also compromised. In order to tackle this issue, gaps in the input sequence \mathbf{P} can be filled with zeros, but the drop in accuracy

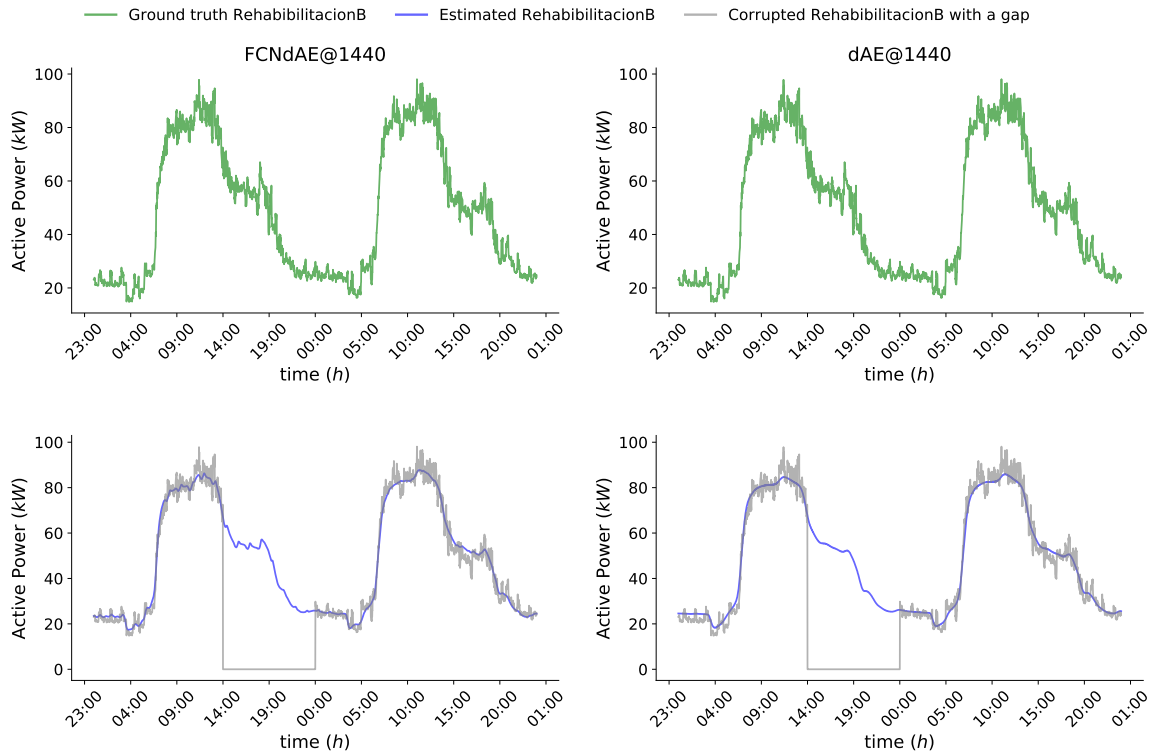


Figure 4.17: Example of gap filling when the failure only affects the individual consumption. This case is exemplified with the estimation of individual node *RehabilitacionB* computed by dAE and FCN-dAE approaches with input sequence length of 1440 minutes.

estimations accuracy may persist, since the amount of input temporal context in D_m is reduced proportionally to the size of the gap.

The effects in accuracy for FCN-dAE and dAE models when gaps in the input sequence \mathbf{P} take place were tested on corrupted test input segments. More specifically, 50 gaps with lengths that range from 0 to 700 minutes have been randomly placed along 1-month test segments of the main consumption.

Figure 4.18 shows the accuracy of the resulting disaggregation for all nodes using RMSE as the evaluation metric. After increasing the gap length, the performance of all approaches decreases, as expected, due to the reduction in input temporal context. However, FCN-dAE models seem to be more stable against large gaps in those individual nodes that previously outperformed dAE models in table 4.1, and therefore they result in more robust approaches for gap filling purposes. This can be explained by their fully-convolutional architecture, where its intermediate activations depend locally on a few elements of the input (receptive field), instead of on the entire input segment, as in the

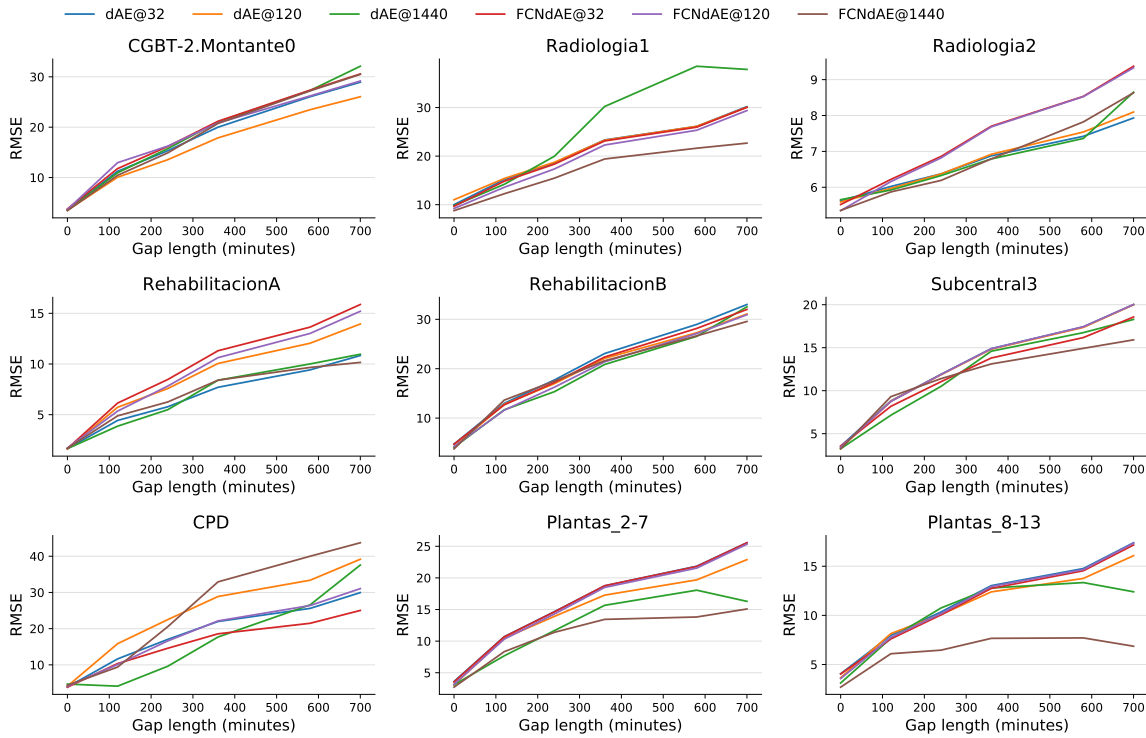


Figure 4.18: Obtained RMSE evolution after applying FCN-dAE and dAE models to corrupted test main consumption segments with different gap lengths.

case of dAE approaches.

Finally, in Figure 4.19, an example of gap filling when 9-hour failures take place in the input sequence is shown. It can be seen how FCN-dAE model seems to need less uncorrupted input elements in order to produce accurate predictions than the dAE approach.

Novelty detection

The studied DNN-based based on *Auto-Encoders* (AE) learn to extract, from the main consumption, estimations of each individual consumption that follow the distribution of the individual consumptions in training data, so that they are not able to generate individual consumption waveforms out from the training data distribution.

Thus, if an abnormal behavior or consumption appears in an individual node, although it would be reflected in the main consumption, dAE-based models will produce individual consumptions without the anomaly. Some works [185, 186] have suggested

CHAPTER 4. NILM APPROACHES FOR LARGE BUILDINGS BASED ON DEEP LEARNING TECHNIQUES

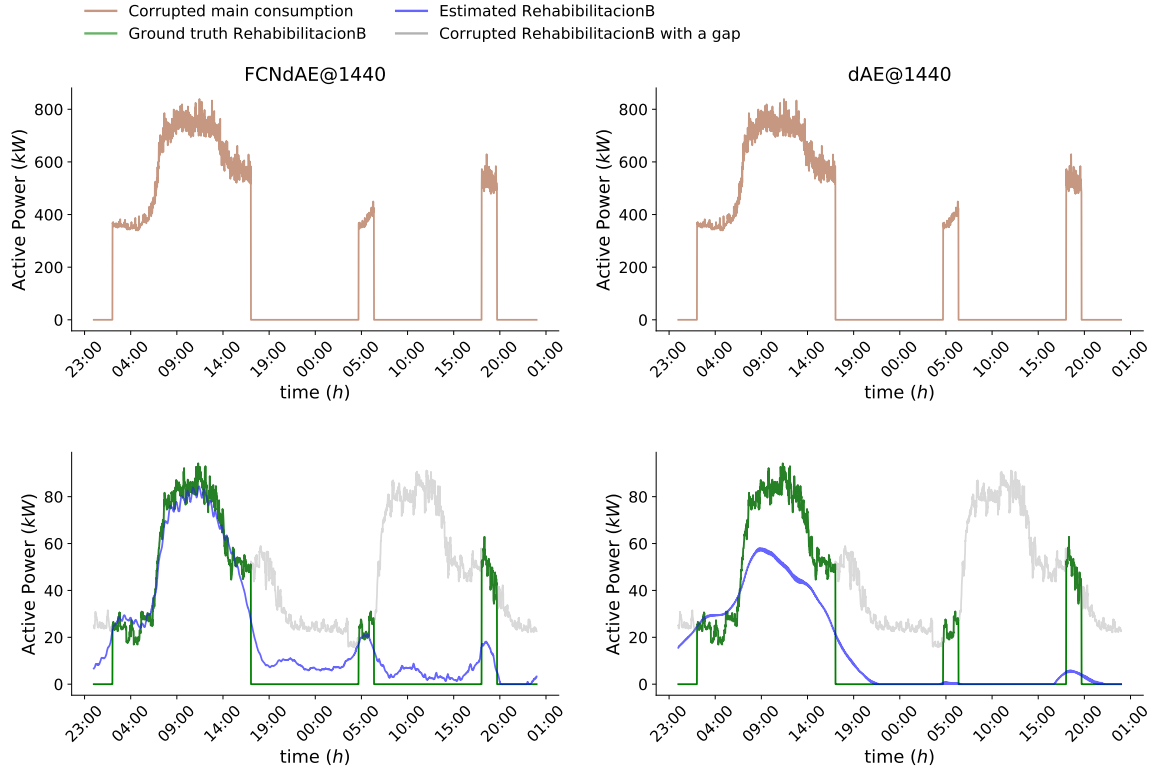


Figure 4.19: Example of gap filling when the failure also affects the main consumption of the facility. This case is exemplified with the estimation of individual node *RehabilitacionB* computed by dAE and FCN-dAE approaches with input sequence length of 1440 minutes and for 9-hour gaps.

that, if the measured real individual consumptions $p_m(t)$ are available, this property can be exploited for detecting potential novelties, using the *residuals* between the real and the estimate output as a novelty score.

In the specific case of NILM models, we can take advantage of the residual between the estimated and the measured individual consumption as novelty score, as it is exemplified in the diagram of Figure 4.20. Hence, the residuals $\text{Res}_m(t)$ of an individual node m are defined as the absolute difference between the estimated and the real consumption for that particular node:

$$\text{Res}_m(t) = |p_m(t) - \hat{p}_m(t)| \quad (4.15)$$

Figure 4.21 shows an example of a novelty synthetically introduced in node *Rehabil-*

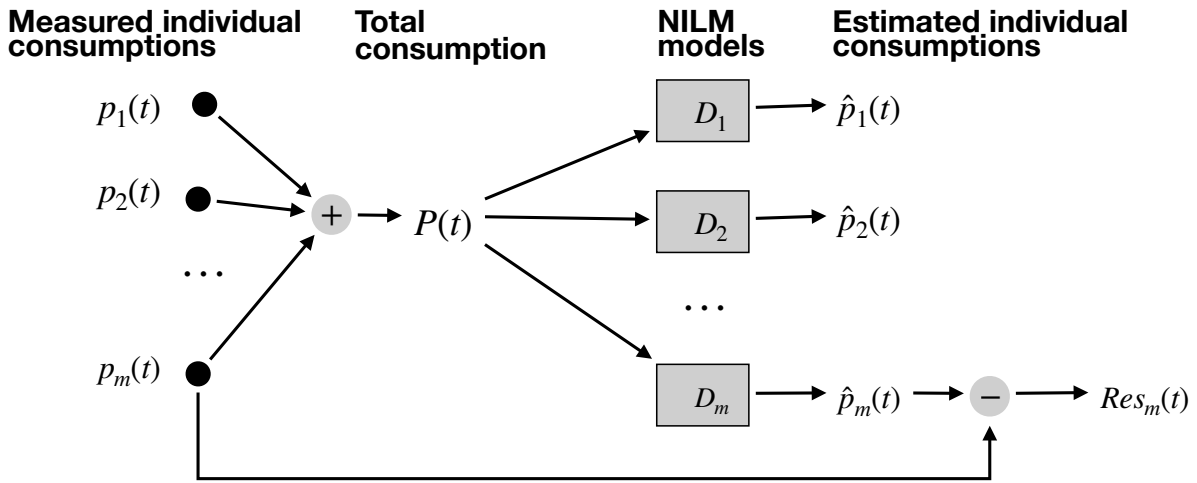


Figure 4.20: Explanatory diagram of the mechanism for obtaining the residuals from the real measured individual consumptions and the estimated individual consumptions obtained.

itacionB, where the residuals clearly highlight the elements of the sequence where the novelty takes place. This means that the FCN-dAE model produces normally distributed estimations, even when main consumption sequences contain important variations. This invariance raises the value of $Res_m(t)$ where modifications take place, showing a significant correlation between the residuals and the behaviors of individual nodes that deviate from the norm.

In order to assess the capacity of $Res_m(t)$ in detecting novelties, trained FCN-dAE and dAE models of section 4.4.2 are applied to a 1-month sequence of main consumption, which has been altered by introducing in its individual nodes the following types of novelties:

1. Gaussian noise with a standard deviation of 20 kW .
2. Clipping the main consumption waveform to a constant value of 20 kW .
3. Adding a sine wave with an amplitude of 20 kW and a period of 100 minutes.

A total of six 6-hours-long novelties per type of novelty were randomly introduced in individual nodes. The modifications are also added to the main consumption, simulating




Figure 4.21: Sequence from node *RehabilitacionB* with a novelty synthetically introduced (red color). In the second row, the residual $|p_m(t) - \hat{p}_m(t)|$ between the estimated and measured individual consumption is shown. In this case the estimation $\hat{p}_m(t)$ is computed by means of a FCN-dAE model with an input sequence length of 1440 minutes.

a real scenario where the main meter records the whole node consumption including the novelty.

After computing the residuals between the modified individual consumption sequences and their corresponding estimations computed by the FCN-dAE with an input sequence length of 1440 minutes, the probability distribution of the residuals for each node of the hospital are shown in Figure 4.22. The probability density of residuals in the abnormal points are biased to the right in most of the individual nodes, which suggests that the residuals may be a good score of novelty.

In light of these results, a binary classification is defined by comparing the residual value $\text{Res}_m(t)$ with a hand-defined threshold θ_m .

$$C_m(t) = \begin{cases} 0, & \text{if } |\text{Res}_m(t)| < \theta_i \\ 1, & \text{otherwise} \end{cases} \quad (4.16)$$

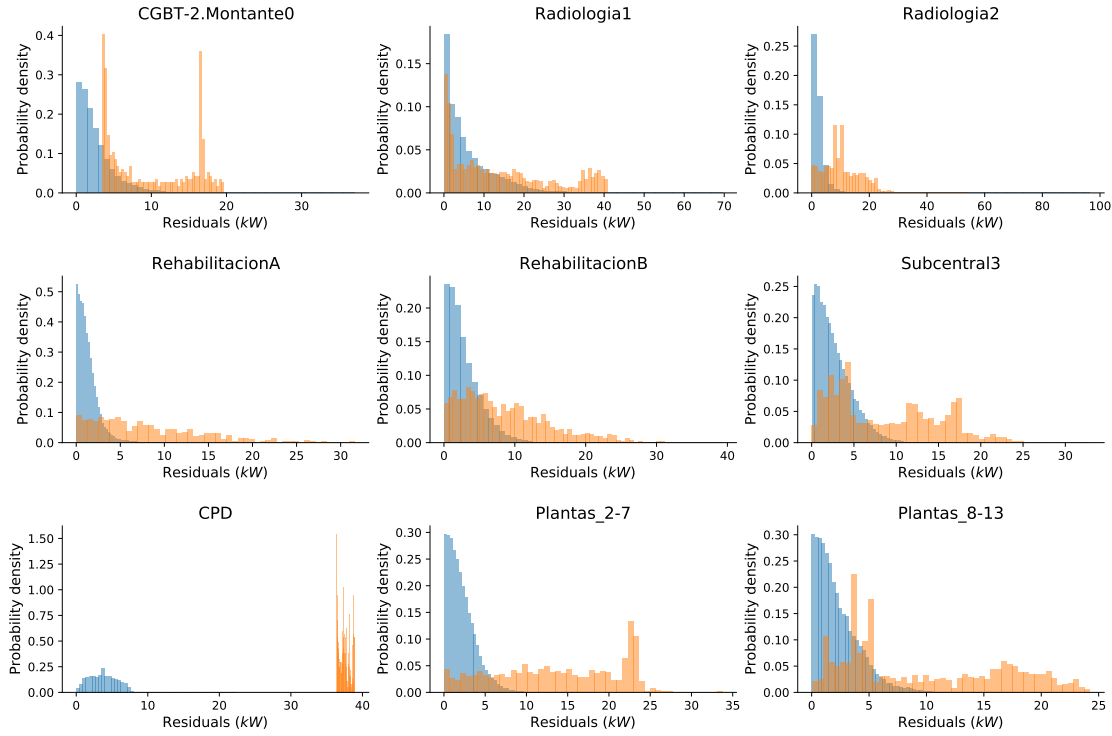


Figure 4.22: Probabilty density of normal and abnormal residuals for all individual nodes computed by means of the estimations produced by the FCN-dAE with 1440 minutes input sequence length.

If the residual value exceeds the threshold, the sample is classified as positive class or novelty. A set of θ_m should be manually defined for each individual node, since the distribution of values of $\text{Res}_m(t)$ in normal scenarios differs from node to node, as it can be seen in Figure 4.22. In Figure 4.23, the precision-recall curves for all nodes are shown to assess the accuracy of C_m in detecting abnormal consumptions for different values of θ_m , when it is applied to the modified sequences of the individual nodes.

In the precision-recall curves, one can see that all models show similar behaviors for different values of θ_m and only in a few of the nodes reach good scores of F1. In addition to that, FCN-dAE models with 1440 minutes input sequences present slightly larger areas under the curve (AUC), reaching F1 scores of 0.7 in several models.

Nevertheless, the difference between models in AUC is minimal, and it could be caused by the nature of the node and the introduced modifications. Despite of these modest results, it is shown that in some of the nodes the novelty detectors of Equation 4.16 can be used to monitor abnormal behaviors without any cost using only the NILM

CHAPTER 4. NILM APPROACHES FOR LARGE BUILDINGS BASED ON DEEP LEARNING TECHNIQUES

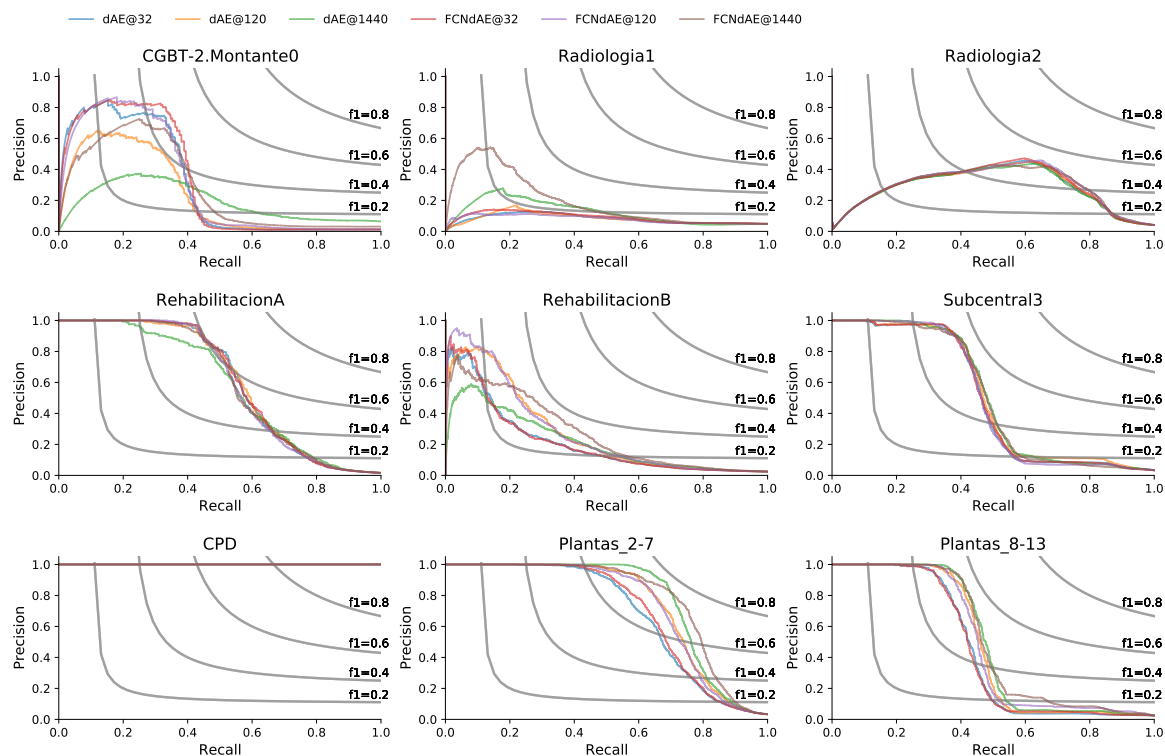


Figure 4.23: Precision-recall curves for residuals computed from all the proposed models ordered by window size. Grey lines indicate the precision-recall curve for a model with the specified F1-score.

model.

4.4.6 Performance of multi-task FCN-dAE

In our multi-task paradigm, the complexity of the overall NILM system is drastically reduced, since all nodes are estimated from a single DNN model instead of training a set of node-specific networks. This reduction of the amount of weights specifically dedicated to each individual consumption may result in a drop in the accuracy of the obtained decompositions.

In order to analyze the effects of the conditioning mechanisms in the accuracy of the disaggregations, the suggested multi-task FCN-dAE is compared with the DNN-based topologies of section 4.4.2. All the models are trained for the Hospital of León dataset, and they are trained with 1440-minutes input sequences. The performance metrics of all these models are gathered from the 1-month test sequence, and they are shown in Table 4.3.

Table 4.3: Comparison of disaggregation performance between sequence-to-sequence models and *multi-task FCN-dAE*.

		biLSTM	dAE	FCNdAE	multi-FCNdAE
MAE	CGBT-2.Montante0	2.9095	2.4863	2.5609	3.2711
	Radiologia1	8.2969	7.0405	5.9346	7.3891
	Radiologia2	3.3974	3.1002	2.7061	3.2705
	RehabilitacionA	1.4721	1.2164	1.2789	1.5251
	RehabilitacionB	3.8967	3.0455	2.8464	3.7287
	Subcentral3	2.7142	2.5553	2.552	3.3734
	CPD	4.2315	4.3056	3.9156	2.6909
	Plantas_2-7	2.4502	2.4512	2.109	3.6736
	Plantas_8-13	2.3686	2.5463	2.0707	4.5218
RMSE	CGBT-2.Montante0	5.2149	3.5161	3.4271	5.318
	Radiologia1	11.6352	9.8098	8.771	10.8085
	Radiologia2	6.0392	5.6539	5.3551	5.7545
	RehabilitacionA	2.1273	1.648	1.6881	2.0971
	RehabilitacionB	6.4429	4.0777	3.7579	5.4655
	Subcentral3	3.5735	3.2491	3.2287	4.195
	CPD	4.6070	4.6968	4.3564	3.1575
	Plantas_2-7	3.1213	3.0926	2.7183	4.5093
	Plantas_8-13	3.0489	3.1191	2.6963	4.9767
NEP	CGBT-2.Montante0	0.0802	0.0685	0.0706	0.0902
	Radiologia1	0.2282	0.1936	0.1632	0.2032
	Radiologia2	0.2749	0.2508	0.2189	0.2646
	RehabilitacionA	0.0701	0.0579	0.0609	0.0726
	RehabilitacionB	0.0978	0.0765	0.0715	0.0936
	Subcentral3	0.1007	0.0948	0.0947	0.1251
	CPD	0.0776	0.0789	0.0718	0.0493
	Plantas_2-7	0.0707	0.0707	0.0609	0.106
	Plantas_8-13	0.1026	0.1103	0.0897	0.1959

As it could be expected, the accuracy of the multi-task FCN-dAE decreases for all individual nodes, except *CPD*. Nevertheless, it is still competitive to the rest of DNN-based models, as it can be seen in the disaggregation shown in Figure 4.24. In the figure, the individual node *Radiologia1* shows that the FCN-dAE model tends to reproduce better the short-term behaviors than its corresponding multi-task version. This is probably caused by the regularization effect of the FiLM layers that may vanish the sharp activations of feature maps, decreasing the capability of the model to reproduce of short-term behaviors at the output. Despite this regularization, the multi-task FCN-dAE model seems to be better in reproducing short-term behaviors of *Radiologia1* than the dAE and the biLSTM models.

Regarding the computational efficiency, when the multi-task FCN-dAE model is

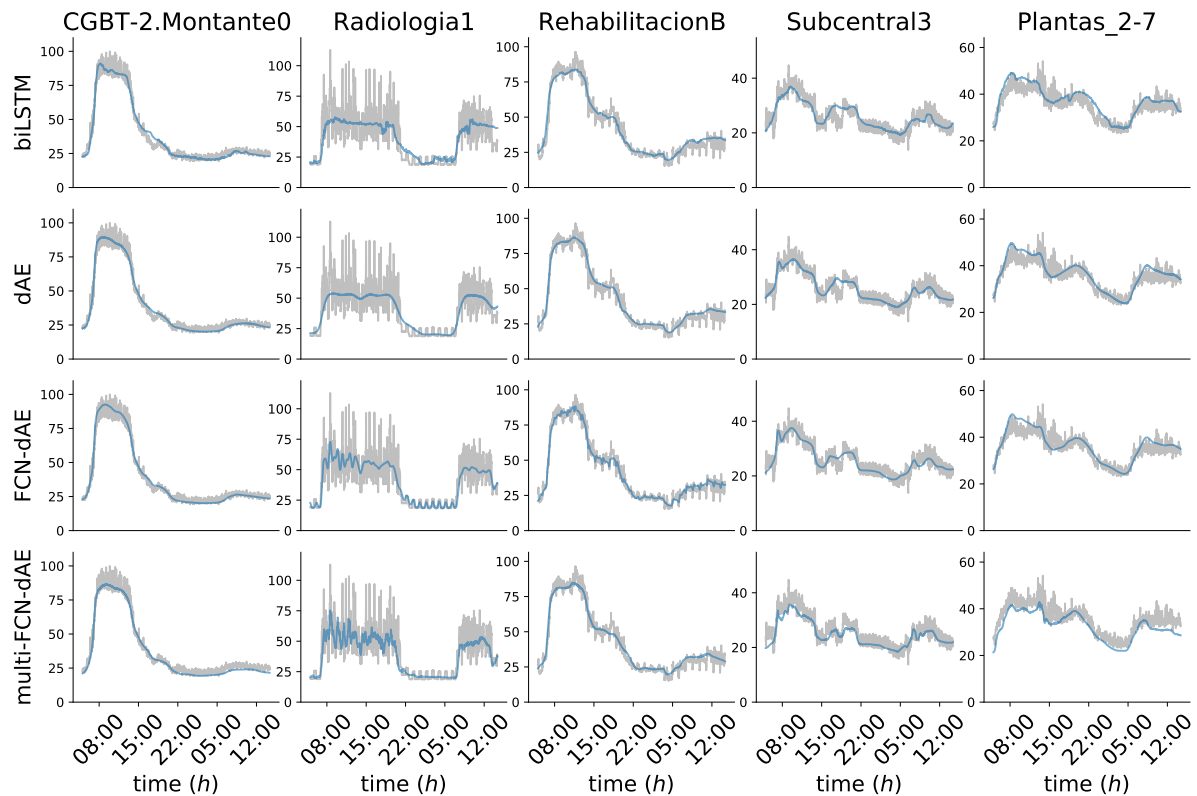


Figure 4.24: Examples of individual nodes from the Hospital of León disaggregated by sequence-to-sequence DNN models and by suggested multi-task FCN-dAE. All models were trained with 1440-minutes input sequences.

compared with previous DNN-based NILM models, it needs more trainable weights than the biLSTM and the vanilla FCN-dAE models due the integration of the FiLM generator, as it is shown in Figure 4.25a. Nevertheless, the suggested multi-task FCN-dAE is able to disaggregate all nodes, so that the total amount of weights needed to estimate all the nodes is much smaller.

The FiLM generator also makes the multi-task FCN-dAE slower than the vanilla FCN-dAE and dAE models in predicting the output individual consumption sequence, as it shown in Figure 4.25a. It takes around 20 milliseconds to process the output sequence, which is still fast enough for our VA paradigm, being close to the limit of a fluid interaction (16 milliseconds) [126].

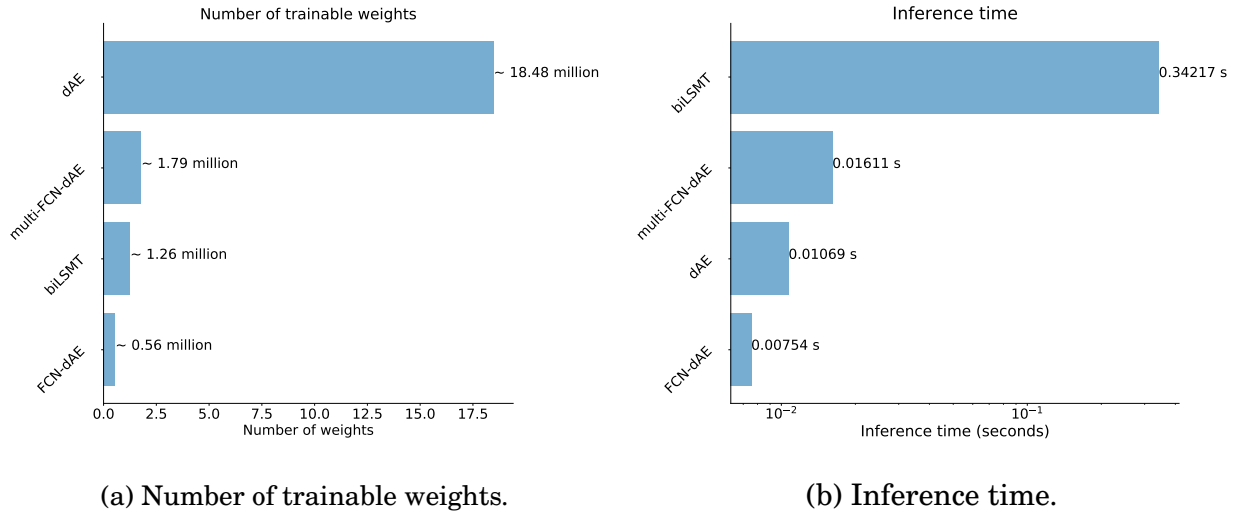


Figure 4.25: Comparison of computational efficiency between suggested multi-task FCN-dAE and sequence-to-sequence DNN-based models.

4.4.7 Interpretation of conditioning mechanisms

The main hypothesis about the FiLM conditioning mechanism incorporated in our multi-task FCN-dAE approach is that it is able to adaptively modulate the encoder feature maps according to the indicated individual consumption, by focusing its attention on those patterns more related to the target consumption. This modulation is not discrete, but in the conditioning input $\mathbf{S}_D^{(i)}$, one can set *continuous transitions* by drawing trajectories between the one-hot encodings, that define the individual nodes in the input space. In Figure 4.26, several examples of continuous transitions between individual nodes are shown.

The smooth transitions between nodes demonstrate the idea that our model continuously *turns off* the features of the starting node and *turns on* the features of the end node. The transitions insightfully reveal to the user, which changes should be applied to the starting individual node to turn it into the end node. This idea is tightly connected to the *explainable machine learning* models based on *counterfactual examples* [187–189], which search the minimum changes to apply on an input sample, so that a DNN-based classifier gives it a different class from the one initially predicted.

The modulation of the feature maps in the encoder also modulates the projection of the input main consumption sequence \mathbf{P} on the FCN-dAE latent space. Therefore, in Figure 4.27, the latent space representations collected from mapping 300 sequences of

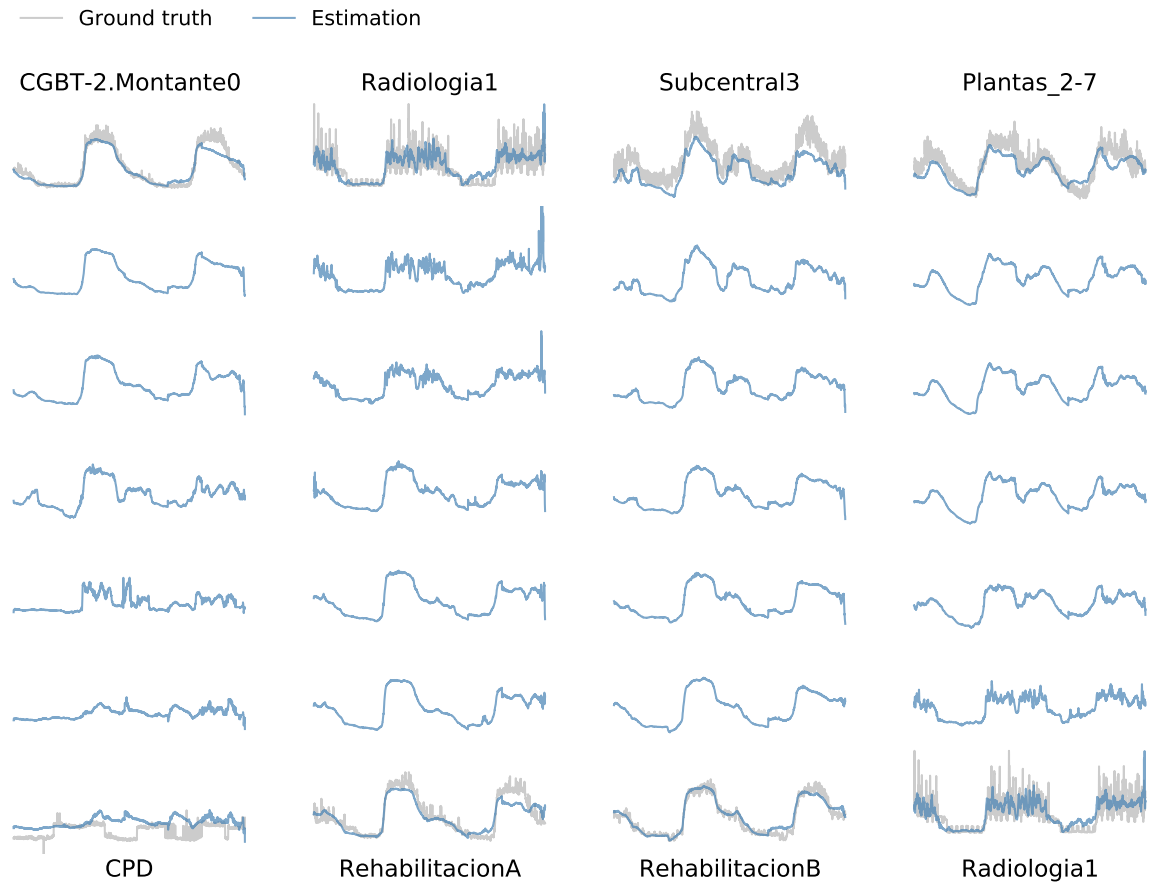


Figure 4.26: Examples of transitions between two individual nodes collected at the output of our multi-task FCN-dAE when continuous trajectories between the one-hot encodings of the nodes are defined in the input S_D .

main consumptions, each of them conditioned several times by all available individual nodes ($300 \text{ sequences} \times 9 \text{ individual nodes} = 2700 \text{ latent space representations}$), are analyzed. All the collected latent space representations are processed by the well-known *Uniform Manifold Approximation and Projection* (UMAP) to visualize a 2D map of the latent space.

The UMAP projections show how the latent space is divided into node-specific subspaces, projecting the same input sequence in different areas according to the indicated individual consumption. Then, the decoder reconstructs each of these subspaces to produce the final individual node estimation. Furthermore, the way in which the latent space is divided seems to be meaningful, since similar consumptions are mapped together, and those consumptions that do not share any similarity, such as the CPD, are placed aside.

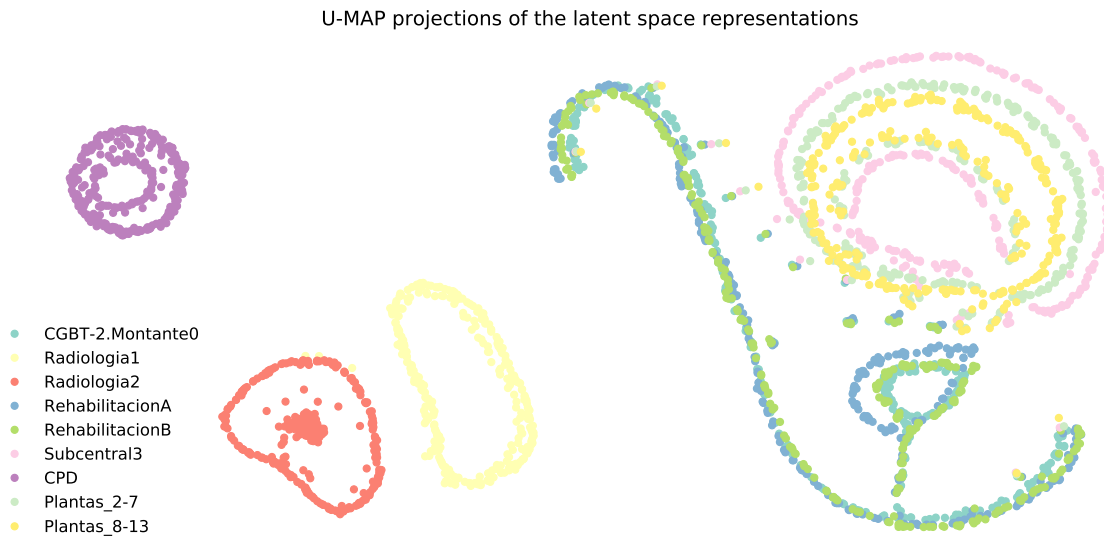


Figure 4.27: UMAP projections of the latent space representation, after apply our multi FCN-dAE model to 300 sequences of main consumptions modulated by all the available individual nodes.

These insights, specially the smooth transitions between nodes, suggest that our multi-task approach provides users with mechanisms to better understand both the inner mechanisms of the DNN model, and the discriminative features of each individual node, than using a set of node-specific FCN-dAEs.

4.4.8 Conclusions

In this chapter, novel DNN techniques have been proposed on the NILM problem in large buildings and compared with DNN-based NILM models previously used in residential facilities. We have identified the input sequence length as a limiting parameter of the DNN-based models applied to large electrical facilities, where the individual consumptions aggregate multiple devices and show a continuous energy demand. A novel fully-convolutional architecture called Fully-Convolutional denoising Auto-Encoder (FCN-dAE) has been proposed as an alternative to vanilla denoising Auto-Encoders (dAE), in this kind of scenarios.

FCN-dAE with a long enough input sequence length have outperformed previous DNN-based models in NILM literature in disaggregating the individual consumptions

from a real Hospital facility. Additionally, FCN-dAE models have shown more capability of disaggregating short and long-term behaviors, as well as more computational efficiency in terms of trainable parameters and inference time, due to its convolutional nature.

In the results section of this chapter, additional potential uses of DNN-based models in the energy management of large facilities are exemplified through several use cases for our hospital dataset. Gap filling, novelty detection or the temporal replacement of real meters are examples of the potential uses of the studied DNN-based NILM systems. In these uses, FCN-dAE models have shown a slightly better performance, specially in filling large gaps in data.

Finally, conditioning mechanisms of the FCN-dAE have been studied with the aim of integrating DNN-based models into the VA paradigm, where the user can modulate the behavior of NILM models according to his intentions, in a fluid manner. Hence, a feature-wise linear modulation has been integrated into our FCN-dAE architecture to provide users with an auxiliary input in which the individual consumption to be disaggregated is indicated, transforming the previous node-specific DNN models into a multi-task model capable of estimating all individual nodes. After testing the resulting multi-task FCN-dAE approach in the hospital data, we have shown that the conditioning mechanisms slightly reduce the accuracy of the resulting disaggregation with respect to vanilla FCN-dAE models, but its performance is comparable to the obtained to state-of-art DNN-based NILM approaches.

Regarding the computational efficiency, the suggested multi-task approach needs less trainable weights than vanilla FCN-dAE models, but it takes more time to process the estimated individual consumption. The delay introduced by the conditioning mechanisms raises the inference time to around the 20 milliseconds in a standard laptop, which is in the limit of a fluid interaction (16 milliseconds), so that it enables the user to reconfigure the disaggregation in a fluid manner.

CONCLUSIONS AND FUTURE WORK

In this last chapter, the final conclusions from the research of this thesis, and the future lines of work are presented. Additionally, the contributions and the main publications resulting from the research carried out during the thesis period are also detailed.

5.1 Discussion and final conclusions

This thesis has addressed the *Non-Intrusive Load Monitoring* (NILM) problem in large buildings from the perspective of *visual analytics*, where users can insightfully understand the whole-installation energy demand through an interactive, intuitive and fluid analysis. The standard NILM framework has been expanded, so that the user can integrate his prior knowledge by means of interaction mechanisms and data visualization techniques.

In chapter 1, our approach is formally introduced by integrating the NILM analysis in the general visual analytics model of Van Wijk [3], and including two novel ways of interaction by which users can steer the behavior of the NILM algorithm and take actions on the facility to improve energy efficiency. The resulting formulation of visual analytics for NILM is not only the root of the subsequent research of this thesis, but also a general purpose model to explain application-centric NILM models that take into account the process, since the user analyzes the NILM outcome, until he or she uses the

acquired knowledge to steer the facility towards a state of better energy efficiency.

To contextualize our research within the NILM literature, in chapter 2, a thorough review of the state of the art in NILM is presented, where we have found that most of the NILM approaches and datasets are suggested for residential facilities, while only a few of them are applied to large non-residential facilities. Hence, the NILM problem in large facilities is identified as an open issue that has to be further studied.

The differences between residential and non-residential facilities in terms of energy demand are explored in the first sections of chapter 3. The analysis of two real and representative datasets of both scenarios reveals that the acquisition of fine-grained individual consumptions in large and complex facilities is unfeasible, and therefore NILM systems in large buildings have to deal with individual consumptions that aggregate multiple devices. This causes that a wide spectrum of events is associated with the same individual consumption, compromising the performance of the previous event-based NILM approaches, and forcing subsequent models to treat the individual consumptions as continuous waveforms.

Taking into account these differences and our visual analytics perspective, an unsupervised NILM approach based on the *Non-negative Matrix Factorization* (NMF) algorithm is suggested in chapter 3. The NMF decomposition has resulted in a highly interpretable, non-negative and sparse decomposition, that can be easily integrated into an efficient *OLAP data cube* structure, paving the way to highly interactive and fluid visualizations. This enables the user to visually correlate the NMF decomposition with context factors, enriching the overall user feedback. Within the mechanisms of interaction, pathways to tweak the NMF algorithm, that allow users to modulate the resulting decomposition are also included.

The idea that the usefulness of a simple NMF decomposition is enhanced by interactive data visualizations is assessed by several use cases in real data from a hospital, where patterns related to the *Heating, Ventilation and Air Conditioning* (HVAC) systems can be easily identified and monitored.

The visual analytics approach in chapter 3 is an example of how the outcome of a simple machine learning algorithm can be extraordinarily enriched by including the user into the analysis. This idea and the use of the data cube to easily incorporate filtering and selection gestures can be transferred to other areas of study, as evidenced by the

participation in other related visual analytics approaches during the research carried out for this PhD thesis. In the analysis of industrial processes, the integration of simple ML algorithms into interactive visualizations boosted by data cube operations have shown to be significantly useful in monitoring and analysis of a cold rolling mill processes, as it is shown in this work [13], in which the author of this PhD thesis has collaborated. In ill-defined problems with a huge number of dimensions, approaches such as the visual analytics approach studied in this thesis can be also applied to spot interesting patterns and find correlations in data throughout a continuous query-answer dialog that broadens the user's knowledge about the problem on each iteration. This is the case of biomedical problems, where fluid reconfigurations of 2D maps based on *morphing projections* [12] have proven to be extremely useful in the exploration of gene expression data.

Although the visual analytics over the NMF decomposition helps to intuitively understand individual consumption patterns, the feedback provided does not correspond to the exact waveforms of individual consumptions. Bearing this in mind, in chapter 4 the use of *deep neural networks* as energy disaggregation models in large buildings is studied. We have introduced a novel DNN-based model called *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE), whose topology composed entirely of convolutional layers facilitates the estimation of short-term and long-term behaviors, outperforming previous state-of-art models in both non-residential and residential data. The convolutional nature of the FCN-dAE is also convenient for our visual analysis perspective due to its computational efficiency, as it reduces the inference time with respect to previous approaches and uses less trainable weights, which opens the door to fluid interaction pathways in subsequent visualizations.

Additionally, in chapter 4 fluid interaction mechanisms between our DNN-based model and the user are studied. We concluded that *conditioning mechanisms* that provide users with elements to modulate the behavior of a neural network are suitable for fast modulations of DNN-based NILM models. More specifically, we have modified our FCN-dAE approach using the *Feature-wise Linear Modulation* (FiLM) approach, which allows us to condition the resulting energy disaggregation on an auxiliary input.

The feasibility of this idea has been proven by defining a *multi-task* model in which the individual consumption to be disaggregated is indicated through the auxiliary input. We have shown that the resulting conditioned approach is able to "select", in a continuous way, an appropriate set of features for the indicated individual node in the auxiliary input,

allowing fluid transitions between the output individual consumptions that improve the interpretation of the estimated consumptions and the internal mechanisms of the model.

The fully-convolutional architectures, together with the conditioning mechanisms exposed in this PhD thesis can be seen as an approach of integrating DNN-based models capable of analyzing time series into visual analytics techniques due to their low computational cost and the interaction with the user. These ideas and the experiments carried out in this work are just a first approach, but this research line has still plenty room of improvement, such as using richer conditioning inputs.

5.2 Thesis contributions

The contributions of the present thesis are summarized as follows:

1. A general *visual analytics* formulation rooted in the model in proposed [3] that makes up for the need of an interpretable analysis of the NILM outcome in order to enhance the energy-saving actions in large facilities has been suggested.
2. We have presented a comparative analysis between the nature of residential and non-residential energy consumptions, by analyzing real and representative datasets in both cases. This analysis has revealed the strong differences between both scenarios, being difficult to measure fine-grained individual consumptions in non-residential facilities, which forces NILM systems to deal with individual consumptions that aggregate multiple devices, and therefore the performance of the previous state-of-art NILM approaches may be compromised, specially those based on sparse events.
3. An unsupervised NILM approach for large buildings, based on our visual analytics NILM formulation, has been proposed. The NILM analysis is tackled as a *blind source separation* problem, where the main consumption of the facility is decomposed into individual components by means of the *non-negative matrix factorization* (NMF) algorithm [5]. The positive nature of the energy consumption is exploited by the NMF decomposition to produce a highly interpretable *parts of a whole* representation, where the user can insightfully visualize and correlate it with context factors of the facility by means of *filter* and *select* interaction mechanisms after arranging the NMF decomposition in an efficient *data cube* structure

[144, 146]. The potential usefulness of this approach has been assessed by several use cases in a real hospital facility, and the approach was firstly presented in the *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* [11], then it was extended to the *Energy and Buildings* journal [10].

4. Taking into account the continuous nature of individual consumptions of large buildings, a novel DNN-based NILM approach completely made of convolutional layers called *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE) has been proposed. We have shown that by replacing all fully-connected layers of previous dAE-based NILM approaches [42] by convolutional layers, following the same paradigm as in *fully-convolutional networks* [170], the resulting network is more robust against variations in the input sequence length, more computationally efficient and capable of reproducing both short-term and long-term behaviors at the output when the input sequence length is large enough. This FCN-dAE approach has been published in the *IEEE Transactions on Smart Grid* journal [9].
5. Several additional potential uses for gap filling and novelty detection of DNN-based models in large facilities have been studied in data from a real hospital, showing the usefulness of NILM systems for daily energy management tasks, and their capability of providing users with valuable feedback.
6. We have integrated the *Feature-wise Linear Modulation* conditioning technique [184] into our FCN-dAE approach, providing users with an interactive mechanism capable of modulating the general behavior of the NILM analysis, and integrating the suggested DNN-based approach in our visual analytics NILM formulation.

5.3 Future work

The ideas and approaches addressed in the present thesis have great potential, and may encourage new research lines that go deeper into some of the studied questions. In this section, several research lines and future work are outlined according to our criteria in order to properly complement the present dissertation.

1. The development of a *user study* for ascertaining the utility of our NMF-based visual analytics solution in the overall understanding of large facilities and in the daily energy management tasks.

2. The improvement of the *interaction pathways between users and the NMF algorithm*, allowing them to introduce constraints into the optimization that modulate the resulting energy decomposition according to their prior knowledge. As an example of such a constraint, authors in [190] have suggested to fix some of the basis consumptions to a known pattern of consumption.
3. New mechanisms of visual interaction that enable to find correlations between the context factors and the outcome of the NILM analysis in order to improve the decision-making process. For example, the *morphing projections* technique [12] provide fast and flexible reconfiguration of 2D views that can be combined with the operations of the data cube in order to create powerful 2D maps, where user can find visual correlations between factors under study.
4. Adding new practices in fully convolutional architectures can be further studied, such as introducing *dilated convolutions* [191], which increase the receptive field of the network without increasing the number of trainable parameters; and integrating *residual connections* [89] that allow to make the networks deeper, since *shortcut connections* mitigate the vanishing and exploding gradient issues.
5. Exploring *new sources of conditioning*, apart from a simple one-hot encoding in our multi-task paradigm. Energy demand profiles or a set of attributes (e.g., weather conditions, occupancy or location of the individual consumption) recognizable for the user that define the individual consumption to be disaggregated are examples of richer sources of conditioning, which enable users to interact with the models by “*asking*” how would the disaggregation be for a user-defined set of attributes (even different for those seen during the training phase). This powerful interaction pathway can help not only to formulate intuitive questions to the model, but also to better understand the internal mechanisms of the trained neural network.
6. The study of the transferability of our conditioned approach to new scenarios, such as the inclusion of new individual nodes or the application of our approach to a new facility. At this regard, the adaptation of the NILM systems may be sped up by leveraging all the learned basic features by the main conditioned network and retraining only the auxiliary network (FiLM generator), that modulates the importance of the features for the new individual consumptions, instead of retraining the entire network, as in previous approaches [111].

CONCLUSIONES Y TRABAJO FUTURO

En este último capítulo se presenten las conclusiones de la investigación realizada en esta tesis doctoral y las líneas de trabajo futuras. Además de ello, las contribuciones y las principales publicaciones de la investigación llevadas a cabo durante el periodo de esta tesis serán también detalladas.

6.1 Discusión y conclusiones finales

Esta tesis ha abordado el problema de la desagregación de energía conocido como *Non-Intrusive Load Monitoring* (NILM) en grandes edificios desde la perspectiva de la *analítica visual*, donde los usuarios pueden entender de una forma intuitiva la demanda energética de toda la instalación mediante un análisis interactivo y fluido. Para ello, el marco de trabajo de NILM ha sido extendido, de manera que el usuario pueda integrar su conocimiento previo sobre la instalación al análisis, mediante mecanismos de interacción y técnicas de visualización de datos.

En el capítulo 1, se ha introducido formalmente nuestro enfoque, donde el análisis NILM es integrado en el esquema general de analítica visual propuesto por Van Wijk [3], introduciéndose dos nuevas vías de interacción mediante las cuales los usuarios pueden guiar el comportamiento del análisis NILM y tomar acciones sobre la instalación con el fin de mejorar la eficiencia energética de la misma. La formulación resultante será nuestro modelo de analítica visual para el problema NILM, el cual no solo será la base para la

posterior investigación, sino que también puede considerarse un modelo de propósito general para explicar modelos NILM centrados en el desarrollo de aplicaciones de usuario que tengan en cuenta todo el proceso desde que el usuario analiza la desagregación obtenida, hasta que él o ella aplique el conocimiento adquirido para guiar la instalación hacia un estado de mejor eficiencia energética.

Para contextualizar nuestra investigación dentro de la literatura de NILM, en el capítulo 2, se revisa en detalle el estado del arte en NILM, donde se concluye que la mayoría de los modelos NILM y los conjuntos de datos disponibles están enfocados a pequeñas instalaciones residenciales, mientras que solo unos pocos son aplicados a grandes instalaciones no residenciales. De este modo, el problema NILM en grandes edificios se identifica como una cuestión abierta que ha de ser estudiada.

Las diferencias entre instalaciones residenciales y no residenciales en términos de demanda energética son exploradas en las primeras secciones del capítulo 3. En ellas, se presenta un análisis exploratorio de dos conjuntos de datos representativos en ambos escenarios, el cual revela que una adquisición detallada de la demanda en los consumos individuales no es viable en grandes instalaciones, y por tanto los sistemas NILM tienen que enfrentarse a consumos individuales que agregan múltiples dispositivos. Esto causa que multitud de eventos en el consumo agregado se asocien al mismo consumo individual, comprometiendo la precisión de los modelos basados en eventos del estado del arte, y forzando a los modelos NILM a tratar los consumos individuales como señales continuas.

Teniendo en cuenta estas diferencias y nuestra perspectiva de analítica visual, en el capítulo 3 se propone un modelo NILM no supervisado basado en el algoritmo *Non-negative Matrix Factorization* (NMF) para grandes edificios. La descomposición NMF proporciona una representación altamente interpretable, positiva y dispersa de la demanda energética que puede ser fácilmente integrada en una estructura de datos tipo *OLAP data cube*, la cual permite ser visualizada e interactuar con ella de forma fluida. La interacción fluida facilita que el usuario establezca correlaciones visuales entre la descomposición NMF resultante y los factores de contexto de la de instalación, enriqueciendo la comprensión general de los resultados. Dentro de los mecanismos de interacción implementados, se incluyen vías de interacción con el propio algoritmo NMF, las cuales permiten a los usuarios modular la desagregación resultante.

La idea de que la utilidad de una simple descomposición NMF puede ser potenciada mediante visualizaciones interactivas es evaluada mediante varios casos de uso reales

con datos procedentes del Hospital Universitario de León, donde se puede comprobar que, gracias a nuestro enfoque de analítica visual, se puede identificar fácilmente patrones relacionados con los sistemas de *climatización* (*Heating, Ventilation and Air Conditioning* HVAC según sus siglas en inglés).

El enfoque de analítica visual presentado en el capítulo 3 es un ejemplo de cómo el resultado de un algoritmo sencillo de aprendizaje automático puede ser extraordinariamente enriquecido incluyendo al usuario en el ciclo de análisis. Esta idea puede ser fácilmente transferida a otras áreas de conocimiento, tal y como se ha demostrado con la participación en otros trabajos en analítica visual durante la investigación llevada a cabo en esta tesis. Dentro del ámbito del análisis de procesos industriales, la integración de algoritmos de aprendizaje automático en visualizaciones interactivas respaldadas con operaciones del cubo de datos han demostrado ser significativamente útil para la monitorización y el estudio de procesos de laminación en frío, tal y como se muestra en [13], donde el autor de la presente tesis ha colaborado activamente. En problemas poco definidos con muchas dimensiones, los enfoques de analítica visual como los abordados en esta tesis pueden ser aplicados para encontrar patrones interesantes en los datos que ayuden al usuario a comprender mejor el problema. Este es el caso de los problemas biomédicos, donde muy a menudo el número de dimensiones es muy alto en comparación con el número de muestras. En estos escenarios se ha demostrado que reconfiguraciones fluidas de mapas 2D de los datos basadas en la técnica *morphing projections* [12] son significativamente útiles en la exploración de conjuntos de datos de expresión genética.

Aunque la analítica visual aplicada sobre la descomposición NMF ayuda a comprender intuitivamente patrones de demanda asociados a subsistemas críticos de los edificios, la información proporcionada no corresponde con la forma de onda exacta de los consumos individuales. Teniendo esto en cuenta, en el capítulo 4, se propone el uso de *redes neuronales profundas* como modelos de desagregación de energía precisos para grandes edificios. Para ello, hemos introducido una novedosa red neuronal completamente convolucional denominada *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE), cuya arquitectura está compuesta enteramente por capas convolucionales, lo que facilita la estimación simultánea de comportamientos a largo y corto plazo en los consumos individuales de salida. La naturaleza convolucional de nuestro enfoque es también apropiada para nuestra formulación de analítica visual, debido a su gran eficiencia computacional, ya que reduce notablemente el tiempo de inferencia y el número de pesos con respecto a modelos basados en redes neuronales previamente utilizados en la literatura NILM.

Además, en el capítulo 4, también se han estudiado técnicas de interacción entre los modelos NILM basados en redes neuronales profundas y el usuario. Hemos concluido que mecanismos de condicionamiento que doten al usuario con elementos para modificar el comportamiento de las redes son apropiados para implementar modulaciones fluidas en los modelos NILM basados en redes neuronales profundas. En concreto, hemos modificado nuestro modelo FCN-dAE, usando la técnica *Feature Linear Modulation* (FiLM), la cual nos permite condicionar la desagregación resultante mediante una entrada auxiliar de la red.

La viabilidad de esta idea ha sido probada definiendo un modelo multitarea, en el cual el consumo individual a desagregar se indica mediante la entrada auxiliar o de condicionamiento. Hemos demostrado que el modelo condicionado resultante es capaz de “seleccionar”, de una forma continua, un conjunto de características apropiadas para el consumo indicado en la entrada condicionando, siendo posible establecer transiciones continuas entre los consumos individuales estimados. Estas transiciones continuas pueden ayudar no solo a la interpretación de los consumos individuales, sino también a comprender los mecanismos internos de las redes neuronales entrenadas.

Las arquitecturas convolucionales junto con los mecanismos de condicionamiento expuestos en esta tesis doctoral pueden ser considerados como una forma preliminar de integrar redes neuronales, capaces de analizar de forma precisa series temporales, en enfoques de analítica visual debido a su bajo coste computacional y la potente interacción con el usuario. Aunque estas ideas y experimentos son solo un primer enfoque, esta línea de investigación tiene un amplio margen de mejora, como la integración de nuevas entradas de condicionamiento que incluyan otro tipo información a parte del consumo individual a desagregar.

6.2 Contribuciones de la tesis

Las contribuciones de esta tesis se resumen en los siguientes puntos:

1. Se ha propuesto una formulación general de *analítica visual* para el problema NILM en grandes edificios basada en el modelo propuesto en [3], que satisface la necesidad de un análisis interpretable del resultado obtenido mediante las técnicas NILM, con el fin de mejorar el proceso de toma de decisiones en materia de ahorro de energía en grandes instalaciones.

2. Hemos presentado un análisis comparativo entre la naturaleza del consumo energético en instalaciones residenciales y no residenciales, mediante la exploración de dos conjuntos de datos representativos en ambos escenarios. Este análisis ha revelado fuertes diferencias entre los dos tipos de instalaciones, siendo más difícil de medir en detalle los consumos individuales en grandes edificios, lo que provoca que los sistemas NILM tengan que enfrentarse a consumos individuales que agregan múltiples dispositivos, y por lo tanto la precisión de los modelos NILM del estado del arte puede verse comprometida, especialmente aquellos basados en eventos.
3. Un modelo NILM no supervisado para grandes edificios, basado en nuestra formulación de analítica visual, ha sido propuesto. El análisis NILM es abordado como un problema de separación de señales, donde el consumo principal de la instalación es descompuesto en componentes individuales mediante el algoritmo *Non-Negative Matrix Factorization* (NMF) [5]. La naturaleza positiva de la demanda energética es explotada por la descomposición NMF para producir una representación basada en partes altamente interpretable, que el usuario puede visualizar intuitivamente y relacionar con los factores de contexto de la instalación mediante mecanismos de filtrado y selección, después de ordenar la descomposición NMF en una estructura tipo *cubo de datos* [144, 146]. La utilidad de este enfoque ha sido evaluada en varios casos de uso con datos reales procedentes de un hospital. Nuestro enfoque NMF fue presentado en la conferencia *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* [11], y el modelo de analítica visual completo, fue enviado y publicado en la revista *Energy and Buildings* [10].
4. Teniendo en cuenta la naturaleza continua de los consumos individuales en grandes edificios, se han adaptado los novedosos enfoques basados en *redes neuronales profundas* a nuestro problema dando lugar a una arquitectura completamente convolucional denominada *Fully-Convolutional denoising Auto-Encoder* (FCN-dAE). Hemos demostrado que reemplazando las capas densas de los modelos dAE anteriores por capas convolucionales, siguiendo el mismo paradigma que en las conocidas *redes fully-convolutional networks* [170], el modelo resultante es más robusto ante variaciones del tamaño de la secuencia de entrada, más eficiente computacionalmente y es capaz de reproducir tanto los comportamientos a largo y corto plazo en la estimación de salida, cuando la ventana de entrada es lo suficientemente grande. El modelo FCN-dAE ha sido recientemente publicado en

la revista *IEEE Transactions on Smart Grid* journal [9].

5. Se han estudiado varios usos potenciales de las técnicas NILM basadas en redes neuronales profundas, tales como el relleno de huecos y la detección de novedades en grandes instalaciones utilizando datos de un hospital real. Esto demuestra la utilidad de los sistemas NILM para las tareas diarias de gestión de la energía, y su capacidad para proporcionar a los usuarios información valiosa acerca del estado de la instalación.
6. Hemos integrado la técnica de condicionamiento *Feature-wise Linear Modulation* en nuestro modelo FCN-dAE, proporcionando a los usuarios un mecanismo de interacción, capaz de modular el comportamiento general del modelo NILM e integrar nuestro modelo FCN-dAE dentro de nuestra formulación de analítica visual en NILM.

6.3 Trabajo futuro

Las ideas y modelos abordados en esta tesis doctoral son, en algunos casos, primeros acercamientos a técnicas con un gran potencial dentro del campo de la eficiencia energética, los cuales pueden sugerir nuevas vías de investigación. En esta sección, de acuerdo a nuestro criterio, se sugieren las siguientes líneas de investigación y trabajo futuro con el fin de complementar la presente disertación:

1. El diseño de un *estudio de usuario* para comprobar la utilidad de nuestro modelo de analítica visual basado en la descomposición NMF, en términos de mejora de la conciencia energética en grandes instalaciones y en las tareas diarias de gestión de la energía.
2. La mejora de las vías de interacción entre los usuarios y el algoritmo NMF, permitiéndoles introducir restricciones en la optimización que modulen la descomposición resultante de acuerdo a su conocimiento experto. Un ejemplo de estas restricciones se sugiere en [190], donde varias bases son excluidas de la optimización NMF e igualadas a consumos o patrones conocidos.
3. Nuevos mecanismos de interacción visual que permitan encontrar correlaciones entre las variables de contexto y la salida del análisis NILM útiles para la mejora de toma de decisiones. Por ejemplo, la técnica *morphing projections* [12], permite

reconfiguraciones fluidas en visualizaciones 2D que pueden ser fácilmente combinadas con las operaciones del cubo de datos para crear mapas 2D, donde el usuario puede encontrar correlaciones visuales entre los factores bajo estudio.

4. Mejoras en la arquitectura FCN-dAE, tales como introducir *dilated convolutions* [191], las cuales incrementan el campo receptivo de la red sin incrementar el número de parámetros a entrenar, o la integración de *conexiones residuales* [89], las cuales permiten crear redes más profundas, ya que mitigan los problemas de desvanecimiento o explosión del gradiente durante el entrenamiento de las mismas.
5. Explorar *nuevas fuentes de condicionamiento* aparte de un simple *one-hot encoding* en nuestro paradigma multitarea. Perfiles de demanda energética o un conjunto de atributos reconocible para el usuario (por ejemplo, condiciones meteorológicas, ocupación del edificio, o la localización del consumo individual), el cual defina el consumo individual a desagregar son ejemplos de fuentes de condicionamiento más ricas en información. Estas nuevas entradas de condicionamiento pueden permitir al usuario interactuar con los modelos “preguntando” como sería la desagregación para un conjunto de atributos definido por él o ella (incluso uno diferente a los utilizados en la fase de entrenamiento). Esta vía de interacción puede ayudar no solo a formular hipótesis más intuitivas, sino también a entender mejor los mecanismos internos de la red.
6. El estudio de la transferibilidad de nuestro modelo NILM condicionado cuando este es aplicado en contextos nuevos, donde, por ejemplo, un nuevo dispositivo es conectado a la red, o cuando nuestro modelo desplegado en una instalación diferente a la utilizada durante la fase de entrenamiento. A este respecto, la adaptación de nuestros modelo FCN-dAE condicionado pueden ser acelerada aprovechando las características aprendidas por la red principal y re-entrenando solo la red auxiliar de condicionamiento (FiLM-generator) que modula la importancia de las características, en vez de re-entrenar la red al completo tal y como se viene haciendo en los modelos transferibles anteriores [111].

AGRADECIMIENTOS

Me gustaría dar las gracias a las siguientes personas e instituciones, sin las cuales la presente tesis no hubiera sido posible:

- Gracias a mi director de tesis, el profesor Ignacio Díaz Blanco, que me dio la oportunidad de iniciar mi carrera como investigador. Me ha inspirado con su ilusión y entusiasmo por el mundo de la investigación, dándome la libertad y la confianza para explorar, probar y discutir nuevas líneas de investigación, ideas y modelos; las inspiradoras discusiones de “pizarra” en el laboratorio son buena prueba de ello.
- El personal del *grupo de investigación GSDPI* [192] y del *Departamento de Ingeniería Eléctrica, Electrónica, Computadores y Sistemas* de la Universidad de Oviedo, que me ha acompañado durante mi periodo de doctorado y me ha ayudado en innumerables ocasiones en el contexto de la investigación y la docencia.
- Al *Hospital Universitario de León* por proporcionar todos los datos utilizados en nuestros experimentos y por todo el conocimiento experto proporcionado para la interpretación de los resultados presentados en este manuscrito.
- Al grupo de investigación *SUPPRESS* [118] de la *Universidad de León* y a su investigador principal, mi co-director, el profesor Manuel Domínguez González. Ellos son los responsables del diseño y mantenimiento del sistema de adquisición de datos del Hospital Universitario de León y quienes me han proporcionado el acceso a los datos y la información necesaria para abordarlos. Agradecerles también su ayuda en los primeros pasos de mi tesis doctoral y su apoyo durante la investigación y las publicaciones realizadas durante mi doctorado.
- Al grupo *Data science group at the Department of Computer and Systems Sciences (DAMI)* de la Universidad de Estocolmo y al profesor Panagiotis Papapetrou. Du-

rante mi breve estancia en DAMI, descubrí el interesante campo del aprendizaje automático interpretable que me inspiró para desarrollar los modelos condicionados sugeridos en el capítulo 4 de mi tesis.

- Al *Gobierno del Principado de Asturias*, el cual, a través de la beca predoctoral “*Severo Ochoa*” (BP16116), me ha proporcionado el apoyo económico necesario para la realización de la presente tesis.
- Al *Ministerio de Economía español (MINECO)* y a los *Fondos FEDER* de la UE bajo la subvención DPI2015-69891-C2-1/2-R. Gracias a su apoyo financiero ha sido posible el desarrollo de los sistemas de adquisición de datos instalados en el hospital y las publicaciones relacionadas con esta tesis.

ACKNOWLEDGMENTS

I would like to thank the following people and institutions:

- My supervisor, Professor Ignacio Díaz Blanco, who gave me the opportunity to start my PhD thesis. He has inspired me with his illusion and enthusiasm for research. He always gave me the freedom and encouraged me to explore, try and discuss new research lines, ideas and models; the inspiring “blackboard” discussions in the lab are proof of this.
- The people from *GSDPI research group* [192] and the *Department of Electric, Electronics, Computers and Systems* from the University of Oviedo, who have accompanied me during my PhD period and have helped me on countless occasions in the context of research and teaching.
- The *Hospital Universitario de León* for providing all the data used in our experiments and for all the expert knowledge provided for the interpretation of the results presented in this manuscript.
- The *SUPPRESS research group* [118] from the *University of León* and his chief researcher and my second supervisor Professor Manuel Domínguez González, who have designed and maintained the data acquisition system of the Hospital Universitario de León. They have guided me through the first steps during my PhD thesis and supported me during my PhD research and publications.
- The *Data science group at the Department of Computer and Systems Sciences (DAMI)* from the University of Stockholm and the Professor Panagiotis Papapetrou. During my short stay in DAMI, I discovered the interesting field of interpretable machine learning which inspired me to develop the conditioned models suggested in chapter 4.

- The *Principado de Asturias government* through the predoctoral grant “*Severo Ochoa*” (BP16116), which give me the financial support to accomplish the present dissertation.
- The *Spanish Ministry of Economy (MINECO)* and the *FEDER funds* from the EU under grant DPI2015-69891-C2-1/2-R. Thanks to their financial support, the development of the systems installed in the hospital and the publications related to this thesis have been possible.



PUBLICATIONS

In this appendix all the publications related to the research presented in this manuscript are enumerated. They are divided into two categories: 1) *principal* publications which include part of the results presented in this manuscript; and 2) *related* publications in which the author of the present research have collaborated.

A.1 Principal publications

2722

IEEE TRANSACTIONS ON SMART GRID, VOL. 12, NO. 3, MAY 2021

Fully-Convolutional Denoising Auto-Encoders for NILM in Large Non-Residential Buildings

Diego García-Pérez¹, Daniel Pérez-López, Ignacio Díaz-Blanco, *Member, IEEE*, Ana González-Muñiz², Manuel Domínguez-González³, and Abel Alberto Cuadrado Vega, *Member, IEEE*

Abstract—Great concern regarding energy efficiency has led the research community to develop approaches which enhance the energy awareness by means of insightful representations. An example of intuitive energy representation is the parts-based representation provided by Non-Intrusive Load Monitoring (NILM) techniques which decompose non-measured individual loads from a single total measurement of the installation, resulting in more detailed information about how the energy is spent along the electrical system. Although there are previous works that have achieved important results on NILM, the majority of the NILM systems were only validated in residential buildings, leaving a niche for the study of energy disaggregation in non-residential buildings, which present a specific behavior. In this article, we suggest a novel fully-convolutional denoising auto-encoder architecture (FCN-dAE) as a convenient NILM system for large non-residential buildings, and it is compared, in terms of particular aspects of large buildings, to previous denoising auto-encoder approaches (dAE) using real electrical consumption from a hospital facility. Furthermore, by means of three use cases, we show that our approach provides extra helpful functionalities for energy management tasks in large buildings, such as meter replacement, gap filling or novelty detection.

Index Terms—Energy efficiency, building energy consumption, NILM, energy disaggregation, denoising auto-encoders.

I. INTRODUCTION

ENERGY disaggregation, also known as *Non-Intrusive Load Monitoring* (NILM), is the computational process of extracting individual consumptions of an electrical installation from a single total measurement. As it was shown in several studies [1]–[4], parts-based representations of the total consumption provide users with insightful feedback, which encourages them to improve their energy savings. This energy awareness makes NILM systems valuable tools that allow users to detect anomalous individual consumptions, obtain per-appliance forecasts and adjust their energy spending plans better.

Manuscript received May 15, 2020; revised September 22, 2020 and December 2, 2020; accepted December 20, 2020. Date of publication December 28, 2020; date of current version April 21, 2021. This work was supported by the Principado de Asturias Government through the Predoctoral Grant “Severo Ochoa.” Paper no. TSG-00737-2020. (*Corresponding author: Diego García-Pérez.*)

Diego García-Pérez, Ignacio Díaz-Blanco, Ana González-Muñiz, and Abel Alberto Cuadrado Vega are with the Electrical Engineering Department, University of Oviedo, 33204 Gijón, Spain (e-mail: diegogarcia@isa.uniovi.es).

Daniel Pérez-López and Manuel Domínguez-González are with the SUPPRESS Research Group, University of León, 24007 León, Spain.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TSG.2020.3047712>.

Digital Object Identifier 10.1109/TSG.2020.3047712

1949-3053 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

Authorized licensed use limited to: UNIVERSIDAD DE OVIEDO. Downloaded on April 22, 2021 at 09:18:06 UTC from IEEE Xplore. Restrictions apply.


In 1992, Hart [5] addressed for the first time the problem of energy disaggregation, defining each connected appliance as a *finite state machine* (FSM) whose transitions are associated with specific electrical signatures or features manually extracted from the total consumption. Hart’s work led to approaches based on *Hidden Markov Models* (HMM) [6], and more specifically on *Factorial Hidden Markov Models* (FHMM) [7]–[12], in which each individual device is modeled by a set of hidden states and the total measured consumption is the sum of the demands associated with the active states per device. Although models based on hidden states achieve accurate results for “on/off” devices, they have several limitations in modeling multi-state appliances and consumptions that are dependent on human behavior. In addition to this drawback, a strong prior knowledge is needed because the feature extraction and the number of states per appliance are still manually defined.

The revival of neural networks in the last decade arose new deep convolutional architectures [13], [14] capable of automatically extracting relevant features for tasks such as computer vision [14]–[16] and speech recognition [17], [18]. These achievements encouraged several works [19], [20] to translate these novel deep neural networks (DNN) approaches into the field of energy disaggregation. Thus, these new perspectives brought novel NILM techniques that outperform the results of previous HMM-based methods using only the raw aggregated active power as input, without manually defining any feature. Subsequent DNN-based works can be grouped in those which continue the paradigm of modeling hidden states, such as previous HMM-based, by means of the use of *recurrent neural networks* (RNN) [21]–[25], and in those which address the NILM as a *denoising problem* [21], [26]–[29]. The latter ones tackle the energy disaggregation as a filtering problem where the total consumption is the noisy signal to be cleaned in order to obtain the individual loads, obtaining remarkable results using filter functions based on *denoising auto-encoders* (dAE) [30].

Most of the previous methods have been tested in residential buildings but only few of them have been applied to large non-residential buildings [31]–[33]. In recent years, interest in energy efficiency of large buildings has increased, since they constitute about 11% of the total energy consumption in developed countries [34]. Therefore, tools as NILM systems, that improve the energy awareness and help users manage the energy demand, could be convenient for

Figure A.1: Article published in the *IEEE Transactions on Smart Grid* journal [9] with the results of chapter 4.

Energy & Buildings 176 (2018) 95–108




ELSEVIER

Contents lists available at [ScienceDirect](#)

Energy & Buildings


journal homepage: www.elsevier.com/locate/enbuild



Interactive visualization for NILM in large buildings using non-negative matrix factorization

Diego García^{a,*}, Ignacio Díaz^a, Daniel Pérez^a, Abel A. Cuadrado^a, Manuel Domínguez^b, Antonio Morán^b

^a Electrical Engineering Department, University of Oviedo, Edif. Departamental, Campus de Viesques s/n, Gijón 33204 Spain
^b SUPPRESS Research Group, University of Leon, Escuela de Ingenierías, Campus de Vegazana, León 24007 Spain



ARTICLE INFO

Article history:
 Received 5 February 2018
 Revised 24 May 2018
 Accepted 28 June 2018
 Available online 20 July 2018

Keywords:
 Energy efficiency
 Visual analytics
 Building energy consumption
 NMF
 Energy disaggregation
 Interactive displays

ABSTRACT

Non-intrusive load monitoring (NILM) techniques have recently attracted much interest, since they allow to obtain latent patterns from power demand data in buildings, revealing useful information to the expert user. Unsupervised methods are specially attractive, since they do not require labeled datasets. Particularly, *non-negative matrix factorization* (NMF) methods decompose a single power demand measurement over a certain time period into a set of components or “parts” that are sparse, non-negative and sum up the original measured quantity. Such components reveal hidden temporal patterns which may be difficult to interpret in complex systems such as large buildings. We suggest to integrate the knowledge of the user into the analysis in order to recognize the real events inside the electric network behind the learnt patterns. In this paper, we integrate the available domain knowledge of the user by means of a visual analytics web application in which an expert user can interact in a fluid way with the NMF outcome through visual approaches such as *barcharts*, *heatmaps* or *calendars*. Our approach is tested with real electric power demand data from a hospital complex, showing how the interpretation of the decomposition is improved by means of interactive *data cube* visualizations, in which the user can insightfully relate the NMF components to characteristic demand patterns of the hospital such as those derived from human activity, as well as to inefficient behaviors of the largest systems in the hospital.

© 2018 The Authors. Published by Elsevier B.V.
 This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

One interesting approach in energy efficiency is the improvement of electric consumption strategies by means of power demand monitoring tools. Thanks to the growing amount of smart meters installed recently, a lot of measurements of power demand are being gathered today in buildings, household and industrial systems, containing useful hidden information that can help in making efficient decisions if it is extracted and presented intuitively. Thus, developing techniques that are able to extract characteristic patterns of how energy is being consumed from large volumes of demand data, as well as methods to visualize these patterns in an efficient and intuitive way have become promising research topics. A suitable approach with these features is *visual analytics* (VA) [1,2]. VA exploits the insightful synergies between *intelligent data analysis* (IDA), *data visualization* and *interaction* mechanisms, allowing the user to get knowledge from efficient visualiza-

tions of raw data and the IDA outcome. By means of interaction mechanisms, the user is brought into the analysis, being able to modify according to his expert knowledge of the system both the visualization and the IDA analysis.

Applying techniques based on IDA algorithms, it is possible to address issues that require a certain learning from the data such as forecasting energy consumption [3–5], getting temporal patterns in electric power demand [6] or the factorization of a total electric power demand into consumptions downstream that have not been measured individually. Decomposing a total consumption, in a sensorless way, is called in the literature *non-intrusive load monitoring* (NILM) [7–9]. This kind of techniques, that can be considered a parts-based representation of total energy, increases the energy awareness of the user, since the obtained components can provide insightful information of how the electric consumptions are distributed temporally and spatially in the network.

Despite recent efforts to develop novel NILM systems, few approaches have explored the NILM outcome within the VA paradigm. NILM techniques are suitable analysis tools for VA, since the obtained disaggregated consumptions reveal temporal patterns

* Corresponding author.
E-mail address: diegogarcia@isa.uniovi.es (D. García).

<https://doi.org/10.1016/j.enbuild.2018.06.058>
 0378-7788/© 2018 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

Figure A.2: Article published in the *Energy and Buildings* journal [10] with the results of chapter 3.

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1. Available from <http://www.i6doc.com/en/>.

Latent variable analysis in hospital electric power demand using non-negative matrix factorization

Diego García¹, Ignacio Díaz¹, Daniel Pérez¹, Abel A. Cuadrado¹, Manuel Domínguez² *

1- Electrical Engineering Dept. University of Oviedo
Edif. Dept., Campus de Viesques s/n 33204, Gijón - SPAIN

2- SUPPRESS Research Group, University of Leon,
Escuela de Ingenierías, Campus de Vegazana, 24007 León, Spain

Abstract. Energy disaggregation techniques have recently attracted much interest, since they allow to obtain latent patterns from power demand data in buildings, revealing useful information to the user. Unsupervised methods are specially attractive, since they do not require labeled datasets. Particularly, *non-negative matrix factorization* (NMF) methods allow to decompose a single power demand measurement over a certain time period into a set of components or “parts” that are sparse, non-negative and sum up the original measured quantity. Such components reveal hidden temporal patterns and events along this period, related to scheduling events and/or demand patterns from subsystems in the network, that are very useful within an energy efficiency context. In this paper we use this approach on demand data from a hospital during a one-year period, using a calendar visualization of the components, revealing relevant facts about the energy expenditure.

1 Introduction

One interesting approach for power demand analysis [1] is *energy disaggregation*, which decomposes an aggregated electrical measurement into several components, revealing features of downstream demand. An example of this can be found in a household, where we are interested in obtaining different appliance consumptions from the overall energy measurement of the house. In recent years, the interest in energy disaggregation or also named *non-intrusive load monitoring* (NILM) has grown, since the perception and knowledge about consumption is improved through parts-based representation of total energy. This kind of representation can also suggest changes in schedules or detect faults in the network. NILM techniques can be classified into supervised and unsupervised methods. In supervised methods, pattern recognition and optimization algorithms [2] are applied, and its principal disadvantage is that a labeled dataset is required. In many cases, obtaining labeled data can increase set-up costs of NILM systems, making unsupervised methods more suitable. Within unsupervised NILM techniques, *blind source separation* (BSS) and *factorial hidden Markov models*

*The authors would like to thank financial support from the Spanish Ministry of Economy (MINECO) and FEDER funds from the EU under grant DPI2015-69891-C2-1/2-R.

Figure A.3: Article published in the *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* [11]. It includes the idea of applying the NMF decomposition to energy-related data.

A.2 Related publications

Bioinformatics, 2020, 1–10
doi: 10.1093/bioinformatics/btaa989
Advance Access Publication Date: 27 November 2020
Original Paper



Data and text mining

Morphing projections: a new visual technique for fast and interactive large-scale analysis of biomedical datasets

Ignacio Díaz^{1,*}, José M. Enguita¹, Ana González¹, Diego García¹, Abel A. Cuadrado¹, María D. Chiara^{2,3} and Nuria Valdés⁴

¹Department of Electrical Engineering, University of Oviedo, Gijón 33204, Spain, ²Institute of Sanitary Research of the Principado de Asturias, Hospital Universitario Central de Asturias, Oviedo 33011, Spain, ³CIBERONC (Network of Biomedical Research in Cancer), Madrid 28029, Spain and ⁴Department of Internal Medicine, Section of Endocrinology and Nutrition, Hospital Universitario de Cabueñes, Gijón 33204, Spain

*To whom correspondence should be addressed.

Associate Editor: Jonathan Wren

Received on May 13, 2020; revised on November 9, 2020; editorial decision on November 13, 2020; accepted on November 16, 2020

Abstract

Motivation: Biomedical research entails analyzing high dimensional records of biomedical features with hundreds or thousands of samples each. This often involves using also complementary clinical metadata, as well as a broad user domain knowledge. Common data analytics software makes use of machine learning algorithms or data visualization tools. However, they are frequently *one-way* analyses, providing little room for the user to reconfigure the steps in light of the observed results. In other cases, reconfigurations involve large latencies, requiring a retraining of algorithms or a large pipeline of actions. The complex and multiway nature of the problem, nonetheless, suggests that user interaction feedback is a key element to boost the cognitive process of analysis, and must be both *broad* and *fluid*.

Results: In this article, we present a technique for biomedical data analytics, based on blending meaningful views in an efficient manner, allowing to provide a natural smooth way to transition among different but complementary representations of data and knowledge. Our hypothesis is that the confluence of diverse complementary information from different domains on a highly interactive interface allows the user to discover relevant relationships or generate new hypotheses to be investigated by other means. We illustrate the potential of this approach with three case studies involving gene expression data and clinical metadata, as representative examples of high dimensional, multidomain, biomedical data.

Availability and implementation: Code and demo app to reproduce the results available at <https://gitlab.com/idiaz/blanco/morphing-projections-demo-and-dataset-preparation>.

Contact: idiaz@uniovi.es

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Biomedical data is growing at astonishing rates with the broadening of access to massive analyses, including laboratory tests, medical image or gene expression data, bearing precise information about the underlying biological state of the subject. However, analysis of high dimensional biomedical datasets is rarely directly managed by doctors, even though they encompass a unique set of challenges, hard important problems and huge potential sanitary impact.

At the same time as quality and availability of biomedical data increased, accordingly did its application in biomedical research.

Most interestingly, there has also been a huge research activity in the field of visualization of this high-dimensional data, as a very powerful tool for researchers, providing better data interpretation, easier detection of patterns and generation of new hypotheses. Related surveys about the use of these techniques can be found in [Kamal *et al.* \(2014\)](#) and [O'Donoghue *et al.* \(2018\)](#).

In this article, we set the focus on the analysis of genomic data for cancer research, which is a paradigmatic example of such problems. Many processes related to the onset, spreading and evolution of cancer, are strongly related to biological pathways that involve complex chains of regulation/deregulation mechanisms acting on

Figure A.4: Article published in the *Oxford Bioinformatics* journal [12]. It is related to the idea of fluid reconfigurations of views according to the user hypotheses similar to the data cube interaction pathways studied in chapter 3.



Contents lists available at ScienceDirect

Computers in Industry

journal homepage: www.elsevier.com/locate/compind

Interactive data visualization of chatter conditions in a cold rolling mill

Daniel Pérez^{a,*}, Ignacio Díaz^a, Abel A. Cuadrado^a, Jose L. Rendueles^b, Diego García^a^aÁrea de Ingeniería de Sistemas y Automática, University of Oviedo, Gijón, Spain
^bArceformittal, Avilés, Spain

ARTICLE INFO

Article history:

Received 6 March 2018
Received in revised form 16 August 2018
Accepted 27 August 2018
Available online 1 October 2018

Keywords:

Visual analytics
Interactive data visualization
Data cube
Rolling mill
Chatter

ABSTRACT

Rolling of flat steel products is an industrial process in the field of metalworking where two or more pairs of rolls reduce the thickness of a steel strip to produce a uniform thickness material. Despite it has been studied for many years, there are still unpredictable problems that can affect the final quality of the product. One of them is the so-called *chatter*, that is a powerful self-excited vibration that appears suddenly and limits the productivity of the process. In this paper, a visual analytics approach is considered for exploratory analysis in order to discover and understand the factors and conditions under which chatter appears. An interactive web-based interface is presented here which allows the user to explore a map of dynamical conditions and visualize relevant details of each chatter onset. A validation case is performed using real data where normal/fault conditions have been identified automatically. By means of interactive exploration, the tool allows to refine an automatic chatter detection method. Moreover, it is shown to reveal correlations between variables, providing in some expected cases data-based confirmation, but also revealing less obvious relationships. Finally, it provides context, allowing to carry out comparative analysis, both qualitative and quantitative, for different subsets of coils (e.g. different years) as well as for different working conditions.

© 2018 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The rolling process transforms the shape of a steel material by means of a thickness reduction passing it between two or more pairs of rolls held by a mill stand. This process is different depending on the temperature of the rolled material. Precisely, a cold rolling mill [21] produces smoother finished products with a uniform exit thickness, commonly performed continuously through several stands in tandem mills. Despite it is an universal process in metalworking, there are still problems that cause economic losses in modern rolling mills. Moreover, the conditions under which these problems arise are not completely understood, making it difficult to prevent their occurrence.

One of the main problems is a self-excited vibration mode called *chatter*, that appears in rolling operations, provoking unacceptable gauge variations in the final surface of the strip, as it is explained in [26]. The removal of chatter is performed by means of a decrease in the rolling velocity. This involves a loss of productivity which makes chatter not only an industrial concern but also an economic one. The analysis of chatter requires understanding the conditions that lead to the instability of the

process. The dynamic interactions between structural phenomena in the mill and the rolling process have been studied along years using theoretical models [18,130,19,31]. However, these models may be too complex with a large amount of parameters to be applied easily or have so many assumptions that excessively simplify the real problem. Tuning these models to suit a particular facility can be very costly requiring the availability of process data and calibration/optimization methods to estimate the model's parameters. On the other hand, rolling mills involve a large amount of interactions featuring coupled thermal, mechanical and computer systems for control. Besides being complex, their dynamics are constantly evolving as a result of changes in the mechanical properties of rolling elements, misadjustments, changes in the working point, etc. Also, the same model needs to be retuned if applied to other mill, even if it is of the same characteristics. All this poses the need for data-based approaches, that are based on the actual behavior of the process.

The current technologies facilitate data acquisition of many parts of a process, described by a large number of variables, and their massive storage in databases is a very common procedure. *Intelligent data analysis* (IDA) algorithms extract information automatically in order to discover new knowledge that may have stayed hidden. A proper visual presentation of the results from these algorithms is an excellent way for communication [27] and efficient interpretation which supports the decision. The so-called

* Corresponding author.
E-mail address: dperez@isa.uniovi.es (D. Pérez).

<https://doi.org/10.1016/j.compind.2018.08.008>

0166-3615/© 2018 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Figure A.5: Article published in the *Computers in Industry* journal [13]. A novel visual analytics solution boosted by a data cube structure similar to the presented in chapter 3 is suggested for the monitoring of a cold roll mill.

BIBLIOGRAPHY

- [1] J. Kelly and W. Knottenbelt, “The uk-dale dataset, domestic appliance-level electricity demand and whole-house demand from five uk homes,” *Scientific data*, vol. 2, no. 1, pp. 1–14, 2015.
- [2] European Environment Agency, “Final energy consumption by sector and fuel in europe.” <https://www.eea.europa.eu/data-and-maps/indicators/final-energy-consumption-by-sector-10/assessment>, 2020.
- [3] J. J. Van Wijk, “The value of visualization,” in *VIS 05. IEEE Visualization, 2005.*, pp. 79–86, IEEE, 2005.
- [4] R. Wirth and J. Hipp, “Crisp-dm: Towards a standard process model for data mining,” in *Proceedings of the 4th international conference on the practical applications of knowledge discovery and data mining*, pp. 29–39, Springer-Verlag London, UK, 2000.
- [5] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [6] T. Munzner, *Visualization analysis and design*. CRC press, 2014.
- [7] P. Pirolli and S. Card, “The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis,” in *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4, McLean, VA, USA, 2005.
- [8] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han, “Unsupervised disaggregation of low frequency power measurements,” in *Proceedings of the 2011 SIAM international conference on data mining*, pp. 747–758, SIAM, 2011.

- [9] D. García-Pérez, D. Pérez-López, I. Díaz-Blanco, A. González-Muñiz, M. Domínguez-González, and A. A. Cuadrado-Vega, “Fully-convolutional denoising auto-encoders for nilm in large non-residential buildings,” *IEEE Transactions on Smart Grid*, 2020.
- [10] D. García, I. Díaz, D. Pérez, A. A. Cuadrado, M. Domínguez, and A. Morán, “Interactive visualization for nilm in large buildings using non-negative matrix factorization,” *Energy and Buildings*, vol. 176, pp. 95–108, 2018.
- [11] D. García Pérez, I. Díaz Blanco, D. Pérez García, A. A. Cuadrado Vega, M. Domínguez González, *et al.*, “Latent variable analysis in hospital electric power demand using non-negative matrix factorization,” in *ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, i6doc. com publication, 2017.
- [12] I. Díaz, J. M. Enguita, A. González, D. García, A. A. Cuadrado, M. D. Chiara, and N. Valdés, “Morphing projections: a new visual technique for fast and interactive large-scale analysis of biomedical datasets,” *Bioinformatics*, 2020.
- [13] D. Pérez, I. Díaz, A. A. Cuadrado, J. L. Rendueles, and D. García, “Interactive data visualization of chatter conditions in a cold rolling mill,” *Computers in Industry*, vol. 103, pp. 86–96, 2018.
- [14] A. McAfee, E. Brynjolfsson, T. H. Davenport, D. Patil, and D. Barton, “Big data: the management revolution,” *Harvard business review*, vol. 90, no. 10, pp. 60–68, 2012.
- [15] T. Mitchell, “Machine learning,” *McCraw Hill*, 1996.
- [16] C. Ware, *Visual thinking for design*. Elsevier, 2010.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- [18] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.

-
- [19] T. Mikolov, A. Deoras, D. Povey, L. Burget, and J. Černocký, “Strategies for training large scale neural network language models,” in *2011 IEEE Workshop on Automatic Speech Recognition & Understanding*, pp. 196–201, IEEE, 2011.
- [20] J. Ma, R. P. Sheridan, A. Liaw, G. E. Dahl, and V. Svetnik, “Deep neural nets as a method for quantitative structure–activity relationships,” *Journal of chemical information and modeling*, vol. 55, no. 2, pp. 263–274, 2015.
- [21] A. Endert, W. Ribarsky, C. Turkay, B. W. Wong, I. Nabney, I. D. Blanco, and F. Rossi, “The state of the art in integrating machine learning into visual analytics,” in *Computer Graphics Forum*, vol. 36, pp. 458–486, Wiley Online Library, 2017.
- [22] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler, “Visual analytics: Scope and challenges,” in *Visual data mining*, pp. 76–90, Springer, 2008.
- [23] C. Turkay, A. Lex, M. Streit, H. Pfister, and H. Hauser, “Characterizing cancer subtypes using dual analysis in caleyo stratomex,” *IEEE computer graphics and applications*, vol. 34, no. 2, pp. 38–47, 2014.
- [24] Y. Liu, E. Jun, Q. Li, and J. Heer, “Latent space cartography: Visual analysis of vector space embeddings,” in *Computer Graphics Forum*, vol. 38, pp. 67–78, Wiley Online Library, 2019.
- [25] K. Ehrhardt-Martinez, K. A. Donnelly, S. Laitner, *et al.*, “Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities,” in *American Council for an Energy-Efficient Economy Washington, DC*, 2010.
- [26] J. Carroll, S. Lyons, and E. Denny, “Reducing household electricity demand through smart metering: The role of improved information about energy saving,” *Energy Economics*, vol. 45, pp. 234–243, 2014.
- [27] E. Wachsmuth, M. Oram, and D. Perrett, “Recognition of objects and their component parts: responses of single units in the temporal cortex of the macaque,” *Cerebral Cortex*, vol. 4, no. 5, pp. 509–522, 1994.
- [28] S. E. Palmer, “Hierarchical structure in perceptual representation,” *Cognitive psychology*, vol. 9, no. 4, pp. 441–474, 1977.

BIBLIOGRAPHY

- [29] G. W. Hart, “Nonintrusive appliance load monitoring,” *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
- [30] Bidgely, “Bidgely company.” <https://www.bidgely.com/technology/>, 2020.
- [31] voltaware, “voltaware company.” <https://voltaware.com/our-technology>, 2020.
- [32] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar, “Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey,” *Sensors*, vol. 12, no. 12, pp. 16838–16866, 2012.
- [33] M. Zeifman and K. Roth, “Nonintrusive appliance load monitoring: Review and outlook,” *IEEE transactions on Consumer Electronics*, vol. 57, no. 1, pp. 76–84, 2011.
- [34] M. Zhuang, M. Shahidehpour, and Z. Li, “An overview of non-intrusive load monitoring: Approaches, business applications, and challenges,” in *2018 International Conference on Power System Technology (POWERCON)*, pp. 4291–4299, IEEE, 2018.
- [35] N. J. Nunes, L. Pereira, F. Quintal, and M. Berges, “Deploying and evaluating the effectiveness of energy eco-feedback through a low-cost nilm solution,” in *Proceedings of the 6th International Conference on Persuasive Technology*, pp. 2–5, 2011.
- [36] S. Barker, S. Kalra, D. Irwin, and P. Shenoy, “Nilm redux: The case for emphasizing applications over accuracy,” in *NILM-2014 workshop*, Citeseer, 2014.
- [37] A. Faustine, N. H. Mvungi, S. Kaijage, and K. Michael, “A survey on non-intrusive load monitoring methodologies and techniques for energy disaggregation problem,” *arXiv preprint arXiv:1703.00785*, 2017.
- [38] C. Klemenjak and P. Goldsborough, “Non-intrusive load monitoring: A review and outlook,” *arXiv preprint arXiv:1610.01191*, 2016.
- [39] J. R. Herrero, Á. L. Murciego, A. L. Barriuso, D. H. de La Iglesia, G. V. González, J. M. C. Rodríguez, and R. Carreira, “Non intrusive load monitoring (nilm): A state of the art,” in *International Conference on Practical Applications of Agents and Multi-Agent Systems*, pp. 125–138, Springer, 2017.

- [40] Z. Wang and G. Zheng, "Residential appliances identification and monitoring by a nonintrusive method," *IEEE transactions on Smart Grid*, vol. 3, no. 1, pp. 80–92, 2011.
- [41] V. Stankovic, J. Liao, and L. Stankovic, "A graph-based signal processing approach for low-rate energy disaggregation," in *2014 IEEE symposium on computational intelligence for engineering solutions (CIES)*, pp. 81–87, IEEE, 2014.
- [42] J. Kelly and W. Knottenbelt, "Neural nilm: Deep neural networks applied to energy disaggregation," in *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, pp. 55–64, 2015.
- [43] M. Berges, E. Goldman, H. S. Matthews, L. Soibelman, and K. Anderson, "User-centered nonintrusive electricity load monitoring for residential buildings," *Journal of computing in civil engineering*, vol. 25, no. 6, pp. 471–480, 2011.
- [44] S. B. Leeb, S. R. Shaw, and J. L. Kirtley, "Transient event detection in spectral envelope estimates for nonintrusive load monitoring," *IEEE Transactions on Power Delivery*, vol. 10, no. 3, pp. 1200–1210, 1995.
- [45] Y. Jin, E. Tebekaemi, M. Berges, and L. Soibelman, "Robust adaptive event detection in non-intrusive load monitoring for energy aware smart facilities," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4340–4343, IEEE, 2011.
- [46] L. K. Norford and S. B. Leeb, "Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms," *Energy and Buildings*, vol. 24, no. 1, pp. 51–64, 1996.
- [47] N. Sadeghianpourhamami, J. Ruysinck, D. Deschrijver, T. Dhaene, and C. Develder, "Comprehensive feature selection for appliance classification in nilm," *Energy and Buildings*, vol. 151, pp. 98–106, 2017.
- [48] L. Farinaccio and R. Zmeureanu, "Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses," *Energy and Buildings*, vol. 30, no. 3, pp. 245–259, 1999.
- [49] M. L. Marceau and R. Zmeureanu, "Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential

- buildings,” *Energy conversion and management*, vol. 41, no. 13, pp. 1389–1403, 2000.
- [50] M. B. Figueiredo, A. De Almeida, and B. Ribeiro, “An experimental study on electrical signature identification of non-intrusive load monitoring (nilm) systems,” in *International Conference on Adaptive and Natural Computing Algorithms*, pp. 31–40, Springer, 2011.
- [51] M. Baranski and J. Voss, “Detecting patterns of appliances from total load data using a dynamic programming approach,” in *Fourth IEEE International Conference on Data Mining (ICDM’04)*, pp. 327–330, IEEE, 2004.
- [52] J. Liang, S. K. Ng, G. Kendall, and J. W. Cheng, “Load signature study—part i: Basic concept, structure, and methodology,” *IEEE transactions on power Delivery*, vol. 25, no. 2, pp. 551–560, 2009.
- [53] A. G. Ruzzelli, C. Nicolas, A. Schoofs, and G. M. O’Hare, “Real-time recognition and profiling of appliances through a single electricity sensor,” in *2010 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pp. 1–9, IEEE, 2010.
- [54] W. Lee, G. Fung, H. Lam, F. Chan, and M. Lucente, “Exploration on load signatures,” in *International conference on electrical Engineering (ICEE)*, vol. 152, 2004.
- [55] H. Y. Lam, G. Fung, and W. Lee, “A novel method to construct taxonomy electrical appliances based on load signaturesof,” *IEEE Transactions on Consumer Electronics*, vol. 53, no. 2, pp. 653–660, 2007.
- [56] H. Najmeddine, K. E. K. Drissi, C. Pasquier, C. Faure, K. Kerroum, A. Diop, T. Jouannet, and M. Michou, “State of art on load monitoring methods,” in *2008 IEEE 2nd International Power and Energy Conference*, pp. 1256–1258, IEEE, 2008.
- [57] A. Cole and A. Albicki, “Nonintrusive identification of electrical loads in a three-phase environment based on harmonic content,” in *Proceedings of the 17th IEEE Instrumentation and Measurement Technology Conference [Cat. No. 00CH37066]*, vol. 1, pp. 24–29, IEEE, 2000.

-
- [58] M. Dong, P. C. Meira, W. Xu, and C. Chung, "Non-intrusive signature extraction for major residential loads," *IEEE Transactions on Smart Grid*, vol. 4, no. 3, pp. 1421–1430, 2013.
- [59] A. I. Cole and A. Albicki, "Data extraction for effective non-intrusive identification of residential power loads," in *IMTC/98 Conference Proceedings. IEEE Instrumentation and Measurement Technology Conference. Where Instrumentation is Going (Cat. No. 98CH36222)*, vol. 2, pp. 812–815, IEEE, 1998.
- [60] H.-H. Chang, "Non-intrusive demand monitoring and load identification for energy management systems based on transient feature analyses," *Energies*, vol. 5, no. 11, pp. 4569–4589, 2012.
- [61] C. Duarte, P. Delmar, K. W. Goossen, K. Barner, and E. Gomez-Luna, "Non-intrusive load monitoring based on switching voltage transients and wavelet transforms," in *2012 Future of Instrumentation International Workshop (FIIW) Proceedings*, pp. 1–4, IEEE, 2012.
- [62] W. Chan, A. T. So, and L. Lai, "Harmonics load signature recognition by wavelets transforms," in *DRPT2000. International Conference on Electric Utility Deregulation and Restructuring and Power Technologies. Proceedings (Cat. No. 00EX382)*, pp. 666–671, IEEE, 2000.
- [63] O. Krystalakos, C. Nalmpantis, and D. Vrakas, "Sliding window approach for online energy disaggregation using artificial neural networks," in *Proceedings of the 10th Hellenic Conference on Artificial Intelligence*, pp. 1–6, 2018.
- [64] P. P. M. do Nascimento, "Applications of deep learning techniques on nilm," *Diss. Universidade Federal do Rio de Janeiro*, 2016.
- [65] C. Zhang, M. Zhong, Z. Wang, N. Goddard, and C. Sutton, "Sequence-to-point learning with neural networks for nonintrusive load monitoring," *arXiv preprint arXiv:1612.09106*, 2016.
- [66] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.

- [67] M. Baranski and J. Voss, "Nonintrusive appliance load monitoring based on an optical sensor," in *2003 IEEE Bologna Power Tech Conference Proceedings*, vol. 4, pp. 8–pp, IEEE, 2003.
- [68] M. Baranski and J. Voss, "Genetic algorithm for pattern detection in nialm systems," in *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583)*, vol. 4, pp. 3462–3468, IEEE, 2004.
- [69] K. Suzuki, S. Inagaki, T. Suzuki, H. Nakamura, and K. Ito, "Nonintrusive appliance load monitoring based on integer programming," in *2008 SICE Annual Conference*, pp. 2742–2747, IEEE, 2008.
- [70] F. M. Wittmann, J. C. López, and M. J. Rider, "Nonintrusive load monitoring algorithm using mixed-integer linear programming," *IEEE Transactions on Consumer Electronics*, vol. 64, no. 2, pp. 180–187, 2018.
- [71] D. Kelly, *Disaggregation of domestic smart meter energy data*. PhD thesis, Imperial College London, 2016.
- [72] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle, "Circuit-level load monitoring for household energy management," *IEEE Pervasive Computing*, vol. 10, no. 1, pp. 40–48, 2010.
- [73] H. Altrabalsi, V. Stankovic, J. Liao, and L. Stankovic, "Low-complexity energy disaggregation using appliance load modelling," *Aims Energy*, vol. 4, no. 1, pp. 884–905, 2016.
- [74] T. Onoda, G. Rätsch, and K.-R. Müller, "Applying support vector machines and boosting to a non-intrusive monitoring system for household electric appliances with inverters," 2000.
- [75] Y.-H. Lin and M.-S. Tsai, "Applications of hierarchical support vector machines for identifying load operation in nonintrusive load monitoring systems," in *2011 9th World Congress on Intelligent Control and Automation*, pp. 688–693, IEEE, 2011.
- [76] M. Figueiredo, A. De Almeida, and B. Ribeiro, "Home electrical signal disaggregation for non-intrusive load monitoring (nilm) systems," *Neurocomputing*, vol. 96, pp. 66–73, 2012.

-
- [77] J. Liao, G. Elafoudi, L. Stankovic, and V. Stankovic, "Non-intrusive appliance load monitoring using low-resolution smart meter data," in *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 535–540, IEEE, 2014.
- [78] S. Giri, M. Bergés, and A. Rowe, "Towards automated appliance recognition using an emf sensor in nilm platforms," *Advanced Engineering Informatics*, vol. 27, no. 4, pp. 477–485, 2013.
- [79] G.-y. Lin, S.-c. Lee, J. Y.-j. Hsu, and W.-r. Jih, "Applying power meters for appliance recognition on the electric panel," in *2010 5th IEEE Conference on Industrial Electronics and Applications*, pp. 2254–2259, IEEE, 2010.
- [80] T. Zia, D. Bruckner, and A. Zaidi, "A hidden markov model based procedure for identifying household electric loads," in *IECON 2011-37th Annual Conference of the IEEE Industrial Electronics Society*, pp. 3218–3223, IEEE, 2011.
- [81] M.-Y. Zhai, "A new graph learning-based signal processing approach for non-intrusive load disaggregation with active power measurements," *Neural Computing and Applications*, vol. 32, no. 10, pp. 5495–5504, 2020.
- [82] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [83] M. Längkvist, L. Karlsson, and A. Loutfi, "A review of unsupervised feature learning and deep learning for time-series modeling," *Pattern Recognition Letters*, vol. 42, pp. 11–24, 2014.
- [84] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [85] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103, 2008.
- [86] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

- [87] C. Olah, “Understanding lstm networks,” 2015.
- [88] K. S. Barsim and B. Yang, “On the feasibility of generic deep disaggregation for single-load extraction,” *arXiv preprint arXiv:1802.02139*, 2018.
- [89] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” 2015.
- [90] K. Chen, Q. Wang, Z. He, K. Chen, J. Hu, and J. He, “Convolutional sequence to sequence non-intrusive load monitoring,” *The Journal of Engineering*, vol. 2018, no. 17, pp. 1860–1864, 2018.
- [91] K. Chen, Y. Zhang, Q. Wang, J. Hu, H. Fan, and J. He, “Scale-and context-aware convolutional non-intrusive load monitoring,” *IEEE Transactions on Power Systems*, vol. 35, no. 3, pp. 2362–2373, 2019.
- [92] M. Kaselimi, E. Protopapadakis, A. Voulodimos, N. Doulamis, and A. Doulamis, “Multi-channel recurrent convolutional neural networks for energy disaggregation,” *IEEE Access*, vol. 7, pp. 81047–81056, 2019.
- [93] T. Wang, T. Ji, and M. Li, “A new approach for supervised power disaggregation by using a denoising autoencoder and recurrent lstm network,” in *2019 IEEE 12th International Symposium on Diagnostics for Electrical Machines, Power Electronics and Drives (SDEMPED)*, pp. 507–512, IEEE, 2019.
- [94] M. Kaselimi, N. Doulamis, A. Doulamis, A. Voulodimos, and E. Protopapadakis, “Bayesian-optimized bidirectional lstm regression model for non-intrusive load monitoring,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2747–2751, IEEE, 2019.
- [95] H. Rafiq, H. Zhang, H. Li, and M. K. Ochani, “Regularized lstm based deep learning model: First step towards real-time non-intrusive load monitoring,” in *2018 IEEE International Conference on Smart Energy Grid Engineering (SEGE)*, pp. 234–239, IEEE, 2018.
- [96] W. He and Y. Chai, “An empirical study on energy disaggregation via deep learning,” in *2016 2nd International Conference on Artificial Intelligence and Industrial Engineering (AIIE 2016)*, Atlantis Press, 2016.

- [97] R. Bonfigli, A. Felicetti, E. Principi, M. Fagiani, S. Squartini, and F. Piazza, “Denosing autoencoders for non-intrusive load monitoring: improvements and comparative evaluation,” *Energy and Buildings*, vol. 158, pp. 1461–1474, 2018.
- [98] M. Valenti, R. Bonfigli, E. Principi, and S. Squartini, “Exploiting the reactive power in deep neural models for non-intrusive load monitoring,” in *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, IEEE, 2018.
- [99] J. Z. Kolter and T. Jaakkola, “Approximate inference in additive factorial hmms with application to energy disaggregation,” in *Artificial intelligence and statistics*, pp. 1472–1482, PMLR, 2012.
- [100] T.-P. Jung, S. Makeig, C. Humphries, T.-W. Lee, M. J. Mckeown, V. Iragui, and T. J. Sejnowski, “Removing electroencephalographic artifacts by blind source separation,” *Psychophysiology*, vol. 37, no. 2, pp. 163–178, 2000.
- [101] J. Kolter, S. Batra, and A. Ng, “Energy disaggregation via discriminative sparse coding,” *Advances in neural information processing systems*, vol. 23, pp. 1153–1161, 2010.
- [102] M. Figueiredo, B. Ribeiro, and A. de Almeida, “Electrical signal source separation via nonnegative tensor factorization using on site measurements in a smart home,” *IEEE Transactions on Instrumentation and Measurement*, vol. 63, no. 2, pp. 364–373, 2013.
- [103] B. Zhao, L. Stankovic, and V. Stankovic, “On a training-less solution for non-intrusive appliance load monitoring using graph signal processing,” *IEEE Access*, vol. 4, pp. 1784–1799, 2016.
- [104] B. Zhao, L. Stankovic, and V. Stankovic, “Blind non-intrusive appliance load monitoring using graph-based signal processing,” in *2015 IEEE global conference on signal and information processing (GlobalSIP)*, pp. 68–72, IEEE, 2015.
- [105] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, “Nilmtk: an open source toolkit for non-intrusive load monitoring,” in *Proceedings of the 5th international conference on Future energy systems*, pp. 265–276, 2014.

- [106] N. Batra, R. Kukuluri, A. Pandey, R. Malakar, R. Kumar, O. Krystalakos, M. Zhong, P. Meira, and O. Parson, “A demonstration of reproducible state-of-the-art energy disaggregation using nilmtk,” in *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 358–359, 2019.
- [107] G. W. Hart, *Prototype nonintrusive appliance load monitor: Progress report 2*. MIT Energy Laboratory, 1985.
- [108] L. Pereira and N. Nunes, “Performance evaluation in non-intrusive load monitoring: Datasets, metrics, and tools—a review,” *Wiley Interdisciplinary Reviews: data mining and knowledge discovery*, vol. 8, no. 6, p. e1265, 2018.
- [109] D. Murray, L. Stankovic, V. Stankovic, S. Lulic, and S. Sladojevic, “Transferability of neural network approaches for low-rate energy disaggregation,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8330–8334, IEEE, 2019.
- [110] S. Makonin and F. Popowich, “Nonintrusive load monitoring (nilm) performance evaluation,” *Energy Efficiency*, vol. 8, no. 4, pp. 809–814, 2015.
- [111] M. Kaselimi, N. Doulamis, A. Voulodimos, E. Protopapadakis, and A. Doulamis, “Context aware energy disaggregation using adaptive bidirectional lstm models,” *IEEE Transactions on Smart Grid*, 2020.
- [112] J. Z. Kolter and M. J. Johnson, “Redd: A public data set for energy disaggregation research,” in *Workshop on data mining applications in sustainability (SIGKDD)*, San Diego, CA, vol. 25, pp. 59–62, 2011.
- [113] M. Zeifman and K. Roth, “Disaggregation of home energy display data using probabilistic approach,” in *2012 IEEE International Conference on Consumer Electronics (ICCE)*, pp. 630–631, IEEE, 2012.
- [114] S. Henriët, U. Şimşekli, B. Fuentes, and G. Richard, “A generative model for non-intrusive load monitoring in commercial buildings,” *Energy and Buildings*, vol. 177, pp. 268–278, 2018.
- [115] N. Batra, O. Parson, M. Berges, A. Singh, and A. Rogers, “A comparison of non-intrusive load monitoring methods for commercial and residential buildings,” *arXiv preprint arXiv:1408.6595*, 2014.

- [116] S. Makonin, B. Ellert, I. V. Bajić, and F. Popowich, “Electricity, water, and natural gas consumption of a residential house in Canada from 2012 to 2014,” *Scientific data*, vol. 3, no. 1, pp. 1–12, 2016.
- [117] D. Murray, L. Stankovic, and V. Stankovic, “An electrical load measurements dataset of United Kingdom households from a two-year longitudinal study,” *Scientific data*, vol. 4, no. 1, pp. 1–12, 2017.
- [118] SUPPRESS research group., “Web page of suppress research group.” <https://suppress.unileon.es/presentation/>, 2020.
- [119] Schneider Electric, “Power Monitoring and Control.” <https://www.se.com/ww/en/product-category/4100-power-monitoring-and-control/?filter=business-2-building-automation-and-control>, Consulted in: 24/12/2020.
- [120] Schneider Electric, “PowerLogic ION7550 meter.” <https://www.se.com/ww/en/product-range/1460-powerlogic-ion7550-%7C-ion7650/>, Consulted in: 24/12/2020.
- [121] Schneider Electric, “PowerLogic PM5000.” <https://www.se.com/ww/en/product/METSEPM5110/pm5110-meter%2C-modbus%2C-up-to-15th-h%2C-1do-33-alarms/>, Consulted in: 24/12/2020.
- [122] Schneider Electric, “PowerLogic PM8000 series.” <https://www.se.com/ww/en/product-range/62252-powerlogic-pm8000-series/>, Consulted in: 24/12/2020.
- [123] Schneider Electric, “PM5560 Meter.” <https://www.se.com/ww/en/product/METSEPM5560/pm5560-meter%2C-2-ethernet%2C-up-to-63th-h%2C-1%2C1m-4di-2do-52-alarms/>, Consulted in: 24/12/2020.
- [124] I. Yarbrough, Q. Sun, D. Reeves, K. Hackman, R. Bennett, and D. Henshel, “Visualizing building energy demand for building peak energy analysis,” *Energy and Buildings*, vol. 91, pp. 10–15, 2015.
- [125] Y. Chen, X. Liang, T. Hong, and X. Luo, “Simulation and visualization of energy-related occupant behavior in office buildings,” in *Building Simulation*, vol. 10, pp. 785–798, Springer, 2017.

- [126] N. Elmqvist, A. V. Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. Jankun-Kelly, “Fluid interaction for information visualization,” *Information Visualization*, vol. 10, no. 4, pp. 327–340, 2011.
- [127] M. N. Schmidt and R. K. Olsson, “Single-channel speech separation using sparse non-negative matrix factorization,” in *Ninth International Conference on Spoken Language Processing*, 2006.
- [128] P. Paatero and U. Tapper, “Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values,” *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994.
- [129] D. D. Lee and H. S. Seung, “Algorithms for non-negative matrix factorization,” in *Advances in neural information processing systems*, pp. 556–562, 2001.
- [130] S. A. Vavasis, “On the complexity of nonnegative matrix factorization,” *SIAM Journal on Optimization*, vol. 20, no. 3, pp. 1364–1377, 2010.
- [131] P. O. Hoyer, “Non-negative matrix factorization with sparseness constraints,” *Journal of machine learning research*, vol. 5, no. Nov, pp. 1457–1469, 2004.
- [132] C.-J. Lin, “Projected gradient methods for nonnegative matrix factorization,” *Neural computation*, vol. 19, no. 10, pp. 2756–2779, 2007.
- [133] C. Boutsidis and E. Gallopoulos, “Svd based initialization: A head start for nonnegative matrix factorization,” *Pattern recognition*, vol. 41, no. 4, pp. 1350–1362, 2008.
- [134] A. Cairo, *The Functional Art: An introduction to information graphics and visualization*.
New Riders, 2012.
- [135] M. Card, *Readings in information visualization: using vision to think*.
Morgan Kaufmann, 1999.
- [136] W. S. Cleveland and R. McGill, “Graphical perception: Theory, experimentation, and application to the development of graphical methods,” *Journal of the American statistical association*, vol. 79, no. 387, pp. 531–554, 1984.

-
- [137] J. Mackinlay, “Automating the design of graphical presentations of relational information,” *Acm Transactions On Graphics (Tog)*, vol. 5, no. 2, pp. 110–141, 1986.
- [138] E. R. Tufte and P. R. Graves-Morris, *The visual display of quantitative information*, vol. 2. Graphics press Cheshire, CT, 1983.
- [139] E. R. Tufte, N. H. Goeler, and R. Benson, *Envisioning information*, vol. 126. Graphics press Cheshire, CT, 1990.
- [140] B. Shneiderman, “The eyes have it: A task by data type taxonomy for information visualizations,” in *The craft of information visualization*, pp. 364–371, Elsevier, 2003.
- [141] D. A. Keim, “Information visualization and visual data mining,” *IEEE transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 1–8, 2002.
- [142] S. Few, *Now you see it: simple visualization techniques for quantitative analysis*. No. Sirsi) i9780970601988, 2009.
- [143] J. S. Yi, Y. ah Kang, J. Stasko, and J. A. Jacko, “Toward a deeper understanding of the role of interaction in information visualization,” *IEEE transactions on visualization and computer graphics*, vol. 13, no. 6, pp. 1224–1231, 2007.
- [144] A. Datta and H. Thomas, “The cube data model: a conceptual model and algebra for on-line analytical processing in data warehouses,” *Decision Support Systems*, vol. 27, no. 3, pp. 289–301, 1999.
- [145] J. Gray, S. Chaudhuri, A. Bosworth, A. Layman, D. Reichart, M. Venkatrao, F. Pellow, and H. Pirahesh, “Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals,” *Data mining and knowledge discovery*, vol. 1, no. 1, pp. 29–53, 1997.
- [146] I. D. Blanco, A. A. C. Vega, D. P. López, M. D. González, S. A. Castro, and M. Á. P. Medrano, “Energy analytics in public buildings using interactive histograms,” *Energy and Buildings*, vol. 134, pp. 94–104, 2017.
- [147] N. Elmqvist and J.-D. Fekete, “Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 3, pp. 439–454, 2009.

BIBLIOGRAPHY

- [148] Z. Liu, B. Jiang, and J. Heer, “immens: Real-time visual querying of big data,” in *Computer Graphics Forum*, vol. 32, pp. 421–430, Wiley Online Library, 2013.
- [149] L. Lins, J. T. Klosowski, and C. Scheidegger, “Nanocubes for real-time exploration of spatiotemporal datasets,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2456–2465, 2013.
- [150] Pandas library, “Pandas library documentation.” https://pandas.pydata.org/pandas-docs/stable/user_guide/10min.html, 2020.
- [151] Scikit-learn library, “Scikit-learn library documentation.” <https://scikit-learn.org/stable/>, 2020.
- [152] Crossfilter, “Crossfilter github repository.” <https://github.com/crossfilter/crossfilter>, 2020.
- [153] Flask python web, “Flask library documentation.” <https://flask.palletsprojects.com/en/1.1.x/>, 2020.
- [154] E. Azar and C. C. Menassa, “Framework to evaluate energy-saving potential from occupancy interventions in typical commercial buildings in the united states,” *Journal of Computing in Civil Engineering*, vol. 28, no. 1, pp. 63–78, 2014.
- [155] J. Chen and C. Ahn, “Assessing occupants’ energy load variation through existing wireless network infrastructure in commercial and educational buildings,” *Energy and Buildings*, vol. 82, pp. 540–549, 2014.
- [156] Z. Ghahramani and M. I. Jordan, “Factorial hidden markov models,” *Machine learning*, vol. 29, no. 2, pp. 245–273, 1997.
- [157] S. Vaseghi, “State duration modelling in hidden markov models,” *Signal processing*, vol. 41, no. 1, pp. 31–41, 1995.
- [158] M. T. Johnson, “Capacity and complexity of hmm duration modeling techniques,” *IEEE signal processing letters*, vol. 12, no. 5, pp. 407–410, 2005.
- [159] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [160] Y. Le Cun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, “Handwritten digit recognition with a back-propagation network,”

- in *Proceedings of the 2nd International Conference on Neural Information Processing Systems*, pp. 396–404, 1989.
- [161] C. Zhang, M. Zhong, Z. Wang, N. Goddard, and C. Sutton, “Sequence-to-point learning with neural networks for non-intrusive load monitoring,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [162] C. Nichiforov, G. Stamatescu, I. Stamatescu, V. Calofir, I. Fagarasan, and S. S. Iliescu, “Deep learning techniques for load forecasting in large commercial buildings,” in *2018 22nd International Conference on System Theory, Control and Computing (ICSTCC)*, pp. 492–497, IEEE, 2018.
- [163] A. Morán, S. Alonso, D. Pérez, M. A. Prada, J. J. Fuertes, and M. Domínguez, “Feature extraction from building submetering networks using deep learning,” *Sensors*, vol. 20, no. 13, p. 3665, 2020.
- [164] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE transactions on neural networks*, vol. 5, no. 2, pp. 157–166, 1994.
- [165] S. Bai, J. Z. Kolter, and V. Koltun, “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” *arXiv preprint arXiv:1803.01271*, 2018.
- [166] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [167] I. Sutskever, J. Martens, and G. E. Hinton, “Generating text with recurrent neural networks,” in *ICML*, 2011.
- [168] A. Graves, “Generating sequences with recurrent neural networks,” *arXiv preprint arXiv:1308.0850*, 2013.
- [169] S. Siami-Namini, N. Tavakoli, and A. S. Namin, “A comparison of arima and lstm in forecasting time series,” in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1394–1401, IEEE, 2018.
- [170] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.

BIBLIOGRAPHY

- [171] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [172] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in *2016 fourth international conference on 3D vision (3DV)*, pp. 565–571, IEEE, 2016.
- [173] V. Dumoulin and F. Visin, “A guide to convolution arithmetic for deep learning,” *arXiv preprint arXiv:1603.07285*, 2016.
- [174] A. Odena, V. Dumoulin, and C. Olah, “Deconvolution and checkerboard artifacts,” *Distill*, 2016.
- [175] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [176] M. Mirza and S. Osindero, “Conditional generative adversarial nets,” *arXiv preprint arXiv:1411.1784*, 2014.
- [177] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *arXiv preprint arXiv:1406.2661*, 2014.
- [178] V. Dumoulin, E. Perez, N. Schucher, F. Strub, H. d. Vries, A. Courville, and Y. Bengio, “Feature-wise transformations,” *Distill*, 2018.
<https://distill.pub/2018/feature-wise-transformations>.
- [179] A. v. d. Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu, “Conditional image generation with pixelcnn decoders,” *arXiv preprint arXiv:1606.05328*, 2016.
- [180] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.
- [181] B. Dhingra, H. Liu, Z. Yang, W. W. Cohen, and R. Salakhutdinov, “Gated-attention readers for text comprehension,” *arXiv preprint arXiv:1606.01549*, 2016.

-
- [182] G. Ghiasi, H. Lee, M. Kudlur, V. Dumoulin, and J. Shlens, “Exploring the structure of a real-time, arbitrary neural artistic stylization network,” *arXiv preprint arXiv:1705.06830*, 2017.
- [183] V. Dumoulin, J. Shlens, and M. Kudlur, “A learned representation for artistic style,” *arXiv preprint arXiv:1610.07629*, 2016.
- [184] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, “Film: Visual reasoning with a general conditioning layer,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [185] J. An and S. Cho, “Variational autoencoder based anomaly detection using reconstruction probability,” *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.
- [186] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei, Y. Feng, *et al.*, “Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications,” in *Proceedings of the 2018 World Wide Web Conference*, pp. 187–196, 2018.
- [187] Y. Goyal, Z. Wu, J. Ernst, D. Batra, D. Parikh, and S. Lee, “Counterfactual visual explanations,” in *International Conference on Machine Learning*, pp. 2376–2384, PMLR, 2019.
- [188] A. Van Looveren and J. Klaise, “Interpretable counterfactual explanations guided by prototypes,” *arXiv preprint arXiv:1907.02584*, 2019.
- [189] I. Karlsson, J. Rebane, P. Papapetrou, and A. Gionis, “Explainable time series tweaking via irreversible and reversible temporal transformations,” in *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 207–216, IEEE, 2018.
- [190] A. Miyasawa, Y. Fujimoto, and Y. Hayashi, “Energy disaggregation based on smart metering data via semi-binary nonnegative matrix factorization,” *Energy and Buildings*, vol. 183, pp. 547–558, 2019.
- [191] N. Kalchbrenner, L. Espeholt, K. Simonyan, A. v. d. Oord, A. Graves, and K. Kavukcuoglu, “Neural machine translation in linear time,” *arXiv preprint arXiv:1610.10099*, 2016.
- [192] GSDPI research group., “Web page of gsdpi research group.” <http://isa.uniovi.es/GSDPI/>, 2020.

