



MODELLING RESIDENTIAL WATER CONSUMERS' BEHAVIOR


From smart metered data
to demand management

Doctoral Dissertation of:
Andrea Cominola

Advisor:
Prof. Andrea Castelletti

Co-Advisor:
PhD. Matteo Giuliani

2016 - XXIX Cycle



Cover page design by Antonio De Bellis



POLITECNICO
MILANO 1863

Politecnico di Milano
Department of Electronics, Information, and Bioengineering
Doctoral Program In Information Technology

**MODELLING RESIDENTIAL WATER
CONSUMERS' BEHAVIOR**

From smart metered data
to demand management

Doctoral Dissertation of:
Andrea Cominola

Advisor:
Prof. Andrea Castelletti

Co-Advisor:
PhD. Matteo Giuliani

Tutor:
Prof. Paolo Bolzern

The Chair of the Doctoral Program:
Prof. Andrea Bonarini

2016 – XXIX Cycle

I chose life over death for myself and my friends...
I believe it is in our nature to explore, to reach out into the unknown.
The only true failure would be not to explore at all.

— Ernest Shackleton

Abstract

Water demand in the residential sector is forecasted to grow in the next decades, under human population growth, urbanization, and climate change. Worldwide experiences have been proving that demand management strategies based on technological, financial, legislative, maintenance, and educational mechanisms can complement supply-side management to meet future demands, potentially leading to significant reductions in residential water consumption, as well as reducing short- and long-term utilities' costs. Yet, the design of effective demand-side management strategies relies on our understanding of consumers' behaviors. Therefore, models that quantitatively describe how water demand is influenced and varies in relation to exogenous uncontrolled drivers, water consumers' and household characteristics, and demand management actions are key to explore water users' response to alternative water demand management strategies, ultimately supporting strategic planning and policy design.

On this regard, the advent of smart meters in the late 1990s made available new water consumption data at very high spatial (household) and temporal (from several minutes up to few seconds) resolution, enabling the development and application of data analytics tools and mathematical models to accurately characterize sub-daily water consumption behaviors, as well as end-use consumption profiles.

The main goal of this thesis is advancing data analysis and mathematical models to extract information on water consumers' behavior out of smart metered data, ultimately informing and proposing recommendations to customized demand management. More specifically, in this thesis we contribute novel methodologies for profiling, analyzing, and modeling residential water consumers based on high temporal and spatial resolution data of residential water consumption, and also coupled with several qualitative and quantitative data describing consumers' psychographic features. The main focus is on water demand modelling and management. However, in this thesis we either assess the inter-portability of the methodologies between the water and energy fields of applications, or present integrated water-energy applications.

The first outcome of this research is the first published comprehensive review of more than 130 studies on high and low resolution residential water demand modelling and management, as well as the development of a common framework for the critical review of the best practices currently existing in this research field. Secondly, we developed two novel Non-Intrusive Load Monitoring algorithms demonstrated to achieve high electric power and water disaggregation performance, to be robust to signal noise and data resolution, and to be portable between the two fields of application of water and energy. Thirdly, we developed two novel descriptive and predictive modelling tools to infer water consumers' habits and routines, as well as identify the most relevant determinants of their water consuming or saving behaviors, at the household level. Finally, we implemented a three-phase data-mining procedure composed of data dimensionality reduction, customer segmentation, and factor mapping to capture heterogeneous water-electricity consumption profiles, highlighting differences between daily time-of-use of water and electricity, and allowing for the characterization of users based on psychographic and behavioral factors. Applications of the developed methods onto synthetic data, as well as onto real-world case studies in Switzerland, Australia, and USA, demonstrated that our tools constitute important progress for an effective and efficient exploitation of smart metering data to develop models of water-energy users' behavior at the household scale, and advance the customization of demand-side management strategies.

Part of this research has appeared (or has to appear) in the following journal publications:

- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2015a. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environmental Modelling & Software* 72, 198–214;
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A. E., 2016. Sparse optimization for automated energy end use disaggregation. *IEEE Transactions on Control Systems Technology* 24 (3), 1044–1051;
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2017. A hybrid signature-based iterative disaggregation algorithm for non-intrusive load monitoring. *Applied Energy* 185, 331–344;
- Cominola, A., Spang, E., Giuliani, M., Castelletti, A., Loge, F., Lund, J., 2016c. Residential water and electricity customer segmentation analysis to inform demand-side management programs. *Environmental Science & Technology* (in preparation).

Sommario

Proiezioni future mostrano che la domanda idrica nel settore residenziale è destinata ad aumentare nei prossimi decenni, sotto la pressione della crescita demografica, dell'urbanizzazione, e dei cambiamenti climatici. Da diverse esperienze progettuali e di ricerca in tutto il mondo è emerso come, al fine di soddisfare le esigenze future, strategie di gestione della domanda fondate su meccanismi tecnologici, finanziari, legislativi, educativi e di manutenzione possono essere complementari ad azioni di gestione della rete di fornitura. L'integrazione tra i due approcci può portare a significative riduzioni del consumo di acqua in contesto residenziale, così come ad una diminuzione dei costi per le società di servizio idrico sia nel breve che nel lungo periodo.

Tuttavia, la progettazione di efficaci strategie di gestione della domanda dipende dal nostro livello di comprensione circa i comportamenti dei consumatori. Per questo motivo, al fine di esplorare come i consumatori d'acqua reagiscono a diverse strategie di gestione della domanda, e in ultimo supportare la pianificazione di queste ultime, appare necessario lo sviluppo di modelli matematici che descrivano in maniera quantitativa come la domanda idrica sia influenzata e vari in relazione a variabili esogene, caratteristiche dei consumatori e delle loro abitazioni, e azioni di gestione della domanda.

A questo proposito, l'avvento dei contatori "smart" alla fine degli anni '90 ha permesso di aumentare la frequenza di monitoraggio dei consumi d'acqua, rendendo così disponibili dati di consumo ad alta risoluzione temporale (da alcuni minuti, fino a pochi secondi) oltre che spaziale (singola abitazione). Questi dati hanno quindi favorito lo sviluppo e l'applicazione di strumenti avanzati di analisi dei dati e modelli matematici per caratterizzare in modo dettagliato i comportamenti di consumo sub-giornalieri, così come i profili di consumo dei singoli usi finali.

All'interno di questo contesto, l'obiettivo principale di questa tesi è fornire un contributo allo sviluppo di tecniche di analisi dei dati e modelli matematici per estrarre informazioni circa il comportamento dei consumatori idrici domestici a partire da dati ad alta frequenza raccolti con contatori smart, così da formulare, in ultima analisi, raccomandazioni per supportare la progettazione

di strategie personalizzate di gestione della domanda.

In particolare, questa tesi contribuisce alla definizione ed implementazione di nuove metodologie per la profilazione, l'analisi e la modellizzazione dei consumatori di acqua ad uso residenziale, sia sulla base di soli dati ad alta risoluzione temporale e spaziale, sia accoppiando questi ultimi con dati che descrivono le caratteristiche socio-psicografiche dei consumatori a livello sia qualitativo che quantitativo. Il fulcro principale della tesi riguarda la modellizzazione e gestione della domanda idrica. In realtà, l'analisi è a più ampio spettro, in quanto parallelamente viene valutata l'inter-portabilità delle metodologie sviluppate tra i settori idrico ed energetico, ovvero sono proposte anche applicazioni congiunte e integrate che considerino simultaneamente i consumi dei due settori.

Il risultato iniziale di questa ricerca è la prima rassegna esaustiva degli studi inerenti la modellizzazione e gestione di domanda idrica domestica con dati ad alta e bassa risoluzione. In questa rassegna proponiamo una framework per la revisione di studi e progetti in tale settore, per poi classificarne ed analizzarne criticamente oltre 130. In secondo luogo, abbiamo sviluppato due nuovi algoritmi per il "Non-Intrusive Load Monitoring", o "disaggregazione" del consumo totale misurato nei vari utilizzi finali. Nella tesi dimostriamo come questi algoritmi siano in grado di raggiungere elevate prestazioni nella disaggregazione di dati di potenza elettrica e consumo d'acqua, sono robusti in caso di rumore sul segnale misurato, e mantengono una buona performance anche in caso di risoluzione dei dati più grossolana, oltre ad essere portabili tra i due campi di applicazione di acqua e energia.

In terzo luogo, abbiamo sviluppato due innovativi strumenti di modellizzazione, rispettivamente il primo descrittivo e il secondo predittivo, per inferire le abitudini e routine di uso d'acqua dei consumatori, nonché individuarne i più rilevanti potenziali fattori causali a livello domestico. Infine, abbiamo implementato una procedura di data-mining composta da tre fasi che si occupano, rispettivamente, di riduzione della dimensionalità dei dati tramite Analisi delle Componenti Principali, segmentazione della clientela (*customer segmentation*), e individuazione dei principali fattori legati a particolari profili giornalieri di consumo orario di acqua e elettricità tramite tecniche di "factor mapping". Questa procedura è in grado di individuare differenze tra i vari profili di consumo in termini di orario di utilizzo di acqua e elettricità, permettendo poi di profilare i singoli utenti in maniera automatica, sulla base dei pattern di consumo stessi e di una serie di fattori psicografici e comportamentali. Tramite diverse applicazioni sperimentali dei metodi proposti, sia su dati generati sinteticamente, sia su dati monitorati in casi di studio reali in Svizzera, Australia e Stati Uniti d'America, dimostriamo come i nostri strumenti permettano un efficace utilizzo dei dati misurati tramite smart meter ai fini dello sviluppo di modelli di comportamento del consumo di acqua ed energia a scala di singola abitazio-

ne, contribuendo così alla progettazione di strategie di gestione della domanda personalizzate sulle diverse tipologie di utente.

Parte della ricerca presentata in questa tesi appare (o apparirà prossimamente) nelle seguenti pubblicazioni scientifiche:

- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2015a. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environmental Modelling & Software* 72, 198–214;
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A. E., 2016. Sparse optimization for automated energy end use disaggregation. *IEEE Transactions on Control Systems Technology* 24 (3), 1044–1051;
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2017. A hybrid signature-based iterative disaggregation algorithm for non-intrusive load monitoring. *Applied Energy* 185, 331–344;
- Cominola, A., Spang, E., Giuliani, M., Castelletti, A., Loge, F., Lund, J., 2016c. Residential water and electricity customer segmentation analysis to inform demand-side management programs. *Environmental Science & Technology* (in preparation).

Acknowledgements

A number of people contributed directly or indirectly to this research and supported me during these years of PhD studies.

First of all, I would like to express my sincere gratitude to my supervisor, Professor Andrea Castelletti. He first arouse my curiosity towards the use of new technologies and models for modelling and managing urban water demand. He enabled and supported this research with inestimable dedication, scientific guidance, trust, and resources. I would like to thank Dr. Matteo Giuliani for supervising this research with Professor Andrea Castelletti. He has been a great mentor for me, and I really appreciate his commitment, his patience, his ability to approach science with extreme clarity, and all the time we spent discussing the various ideas and technicalities behind this research.

I am thankful to Professors Jay Lund and Edward Spang for giving me the opportunity of further developing this research during my visiting period at the University of California, Davis. They provided me with new ideas and resources to develop an integrated study on residential water-electricity behaviors, presented in the last chapter of this manuscript.

The research leading to these results has received partial funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 619172 (SmartH2O: an ICT Platform to leverage on Social Computing for the efficient management of Water Consumption). My gratitude goes to all the members of the consortium of SmartH2O. In particular, I wish to thank Professor Andrea Emilio Rizzoli, University of Applied Sciences of Southern Switzerland, and Dr. Dario Piga, IMT School for Advanced Studies Lucca, for the joint research we have been developing on water and electricity end-use disaggregation, and Professor Piero Fraternali, Politecnico di Milano, for encouraging and pushing the development and future exploitation of SmartH2O.

I wish to express my gratitude to Professor David Rosenberg and Adel Abdallah, Utah State University. Their support and assistance has been essential for the development and improvement of part of this work. Also, I appreciate the support of my coauthors Professor Martin Anda, Murdoch University, and Professor Frank Loge, UC Davis.

I would like to acknowledge Professor Rodney Stewart, Griffith University, and Professor Dragan Savic, University of Exeter, for reviewing my PhD thesis and for their constructive comments and suggestions. Their inputs have been essential to add value to the finalized version of this thesis.

I am full of admiration for all the members of the Natural Resources Management group at Politecnico di Milano and my PhD fellows at the Department of Electronics, Information, and Bioengineering. Our interaction and friendship enriched me as a person and a professional. They made my time here pleasant, and knowing that young researchers like them are out there makes me feel hopeful. I would like to thank also the students I co-supervised during these years and who contributed to this research.

Lastly, I would like to express my deepest gratitude to those for whom words are never enough. Francesca, my family, my friends.

Contents

List of Figures	XIII
List of Tables	XVII
Acronyms	XIX
1 Introduction	1
1.1 Residential water demand management	1
1.2 Residential water-energy nexus	3
1.3 Opportunities and challenges for modelling residential water (and energy) demand	5
1.4 Thesis motivation, objectives, and scope	7
1.5 Research Method	12
1.6 Thesis outline	13
1.6.1 Chapter 2	13
1.6.2 Chapter 3	13
1.6.3 Chapter 4	14
1.6.4 Chapter 5	14
1.6.5 Chapter 6	15
2 Using smart meters for advancing residential water demand model- ing and management: state-of-the-art review	17
Abstract	17
2.1 Introduction	18
2.2 Data gathering	21
2.2.1 Low resolution data	21
2.2.2 High resolution data	23
2.2.3 Research challenges and future directions	24
2.3 Water end-uses characterization	26
2.3.1 Trace Wizard	26
2.3.2 Identiflow	27
2.3.3 HydroSense	28

2.3.4	SEQREUS project	28
2.3.5	Research challenges and future directions	29
2.4	User modeling	31
2.4.1	Descriptive models	31
2.4.2	Predictive models	32
2.4.3	Research challenges and future directions	35
2.5	Personalized water demand management strategies	36
2.5.1	Research challenges and future directions	38
2.6	Discussion	40
3	Water and energy end-use disaggregation	43
	Abstract	43
3.1	Introduction	44
3.2	Problem formulation and related work	48
3.3	Methods	51
3.3.1	Sparse optimization-based algorithm	51
3.3.1.1	Theoretical concepts	51
3.3.1.2	Practical implementation	55
3.3.2	Hybrid Signature-based Iterative Disaggregation algorithm	56
3.3.2.1	Supervised Hybrid Signature-based Iterative Dis-	
	aggregation (HSID) algorithm	57
3.3.2.2	Semi-supervised HSID algorithm	63
3.4	Applications to power load end-use disaggregation	63
3.4.1	Power load disaggregation through Sparse Optimization	
	Disaggregation (SOD)	64
3.4.1.1	Experimental settings	64
3.4.1.2	Numerical results	67
3.4.2	Power load disaggregation through HSID	70
3.4.2.1	Experimental settings	70
3.4.2.2	Numerical results	73
3.4.3	Multi-resolution power load disaggregation through su-	
	pervised HSID	83
3.4.3.1	Experimental settings	83
3.4.3.2	Numerical results	83
3.5	Water end-use disaggregation through HSID	86
3.5.1	Developing a stochastic simulation model for the gener-	
	ation of residential water end-use demand time series . .	87
3.5.1.1	Method	87
3.5.2	Multi-resolution water disaggregation through supervised	
	HSID	91
3.5.2.1	Experimental settings	91
3.5.2.2	Numerical results	92

3.6	Discussion	94
4	Descriptive and predictive modelling of residential water users	97
	Abstract	97
4.1	Introduction	98
4.2	Profiling residential water users' routines by eigenbehavior modelling	100
4.2.1	Methodology	100
4.2.1.1	Data pre-processing	100
4.2.1.2	Eigenbehavior extraction through Principal Component Analysis	101
4.2.1.3	Eigenbehavior clustering through k-means	101
4.2.2	Case study	102
4.2.3	Numerical results	102
4.3	Modelling residential water consumers' behaviors by feature selection and feature weighting	106
4.3.1	Methodology	106
4.3.1.1	Feature extraction	107
4.3.1.2	Model learning	109
4.3.2	Case study	110
4.3.2.1	The H2ome Smart project	110
4.3.2.2	Data pre-processing	110
4.3.3	Numerical results	112
4.3.3.1	Feature selection and feature weighting	112
4.3.3.2	Interpretation of the feature extraction results	114
4.3.4	Model learning and users' water consumption profile forecasting	117
4.4	Discussion	120
5	Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs	123
	Abstract	123
5.1	Introduction	124
5.2	Methodology	126
5.2.1	Eigenbehavior extraction	126
5.2.2	Consumer segmentation	128
5.2.3	Factor mapping	129
5.3	Case study	129
5.4	Numerical results	130
5.4.1	Analysis of coupled daily water-electricity consumption profiles	130

Contents

- 5.4.2 Recommendations for the design of customized demand-side management strategies 138
 - 5.4.2.1 Water-Energy conservation programs 138
 - 5.4.2.2 Water peak shifting programs 140
- 5.5 Discussion 141
- 6 Conclusions and future research 145**
- Bibliography 149**

List of Figures

1.1	Simplified framework reporting the phases of residential water demand modelling and management studies. The objectives of core chapters are briefly reported in each box.	9
1.2	Overview of the research method adopted in this thesis.	11
2.1	Five-years count of the 134 publications reviewed in this study.	19
2.2	Flowchart of the general procedure for studying residential water demand management.	22
3.1	Typical power load signatures for five indoor appliances.	57
3.2	Flowchart of the HSID algorithm. Boxes enumeration is consistent with the one adopted in Section 3.3.2.	58
3.3	Validation dataset D_T^V : aggregate electric power consumption from June 1, 2012 to June 30, 2012	66
3.4	Disaggregate power consumption profiles. Results obtained through the optimization-based approach presented in the paper (left panels) and through the Factorial Hidden Markov Model (FHMM)-based approach (right panels).	68
3.5	F-score and R2 disaggregation accuracy on four appliances with increasing number of FHMM states.	74
3.6	Disaggregation performance metrics for supervised HSID and 2-state FHMM algorithm on 5 simultaneously operating appliances. From top: F-score, assigned Power Contribution Error (PCE) and R2 score. Data refer to the Summer period. The performance on the Winter period were found to be very similar.	75
3.7	Consumption patterns estimated at the end use by the supervised HSID algorithm, compared with observed values. A stacked area with all the overlapping end-uses is represented in the top part of the figure. The four charts in the bottom part represent a detailed comparison of model output vs observed values of the 5 considered end-uses.	76

3.8	Disaggregation performance metrics for HSID under different levels of noise on the smart-metered trace to disaggregate. From top: F-score, PCE and R2 score.	79
3.9	Disaggregation performance metrics for semi-supervised HSID. From top: F-score, assigned PCE and R2 score.	82
3.10	Heat pump groundtruth trajectory (black line) and trajectory estimated by HSID (red line) at different resolutions.	85
3.11	Sample signatures for standard-efficiency toilet and shower fixtures.	88
3.12	Samples of end-use trajectories output of the stochastic model for a single household over few hours (left panel) and one day (right panel). For the sake of visualization, the data are aggregated at one minute resolution (left panel) and one hour resolution (right panel).	91
3.13	Sample trajectory of 1-minute total water consumption data, generated for a 2-occupants house.	94
4.1	Empirical distribution of nonzero water consumption readings, with the shading representing the three selected classes L1, L2, L3.	103
4.2	Partitions of hourly consumption reading data across the four consumption classes L0, L1, L2, L3.	104
4.3	Results of users' segmentation with respect to the values of the first eigenbehaviors. Each column is the first eigenbehavior of a user and each row corresponds to the hour of the day for each consumption class L0, L1, L2, L3. Color scale represents the value of eigenbehaviors coefficients.	105
4.4	Relative frequency of average hourly consumption for each profile in each consumption class. The gray area represents standard deviation.	106
4.5	Selection frequency of candidate features over multiple feature selection algorithms runs. Considered predictant: Y_{daily} (left) and Y_{pcDaily} (right).	113
4.6	Weighting ranking of candidate features over multiple feature selection algorithms runs. Considered predictant: Y_{daily} (left) and Y_{pcDaily} (right).	114
4.7	Median daily water consumption and median per-capita daily water consumption for different numbers of house occupants.	115
4.8	Median daily water consumption and median per-capita daily water consumption for houses with different number of toilets.	116
4.9	Median daily water consumption and median garden area for households adopting different irrigation techniques.	116

4.10	Median daily water consumption and median per-capita daily water consumption for different types of house.	117
4.11	Predictive models accuracy based on Feature Weighting (FW) algorithms outputs. The following FW algorithm predictant are represented: Y_{daily} and Information gain FW (top-left), Y_{pcDaily} and Information gain FW (top-right), Y_{daily} and Chi-square score FW (bottom-left), Y_{pcDaily} and Chi-square score FW (bottom-right).	119
5.1	Flowchart of the methodology.	127
5.2	Histogram representing the distribution of variance explained [%] by the first eigenbehavior of each user.	131
5.3	Water-and-energy eigenbehavior coefficients and clusters. Each column of the matrix represents the coefficients of the first eigenbehavior for one user (each user is identified by means of a user ID, see x-axis). Eigenbehavior coefficients are reported for diverse labelled levels of water and energy consumption over 24 hours (see y-axis). Color intensity is proportional to coefficient magnitude, as represented in the colorbar on the right side. Eigenbehaviors are clustered in 18 groups based on their similarity: clusters are identified by the labels P1, P2, ..., P18 on top of the figure and the vertical dashed lines.	132
5.4	Empirical Cumulative Distribution Function of hourly water consumption data [cfh], for each profile identified in the clustering phase.	134
5.5	Empirical Cumulative Distribution Function of hourly electricity consumption data [kWh], for each profile identified in the clustering phase.	134
5.6	Median weights of the first eigenbehavior for each profile P1, P2, ..., P18. Profiles are sorted in decreasing order of total per-capita water consumption during the observation period (from top-left to bottom-right). Solid lines refer to water, dashed lines to electricity. Colors refer to different consumption labels.	137
5.7	Average per-household daily water and electricity consumption for each consumer cluster. Each scatter point represents a cluster of consumers. Clusters falling inside the dark and light blue boxes are potential targets for water and energy conservation interventions.	139

List of Tables

2.1	Details of the papers reviewed.	19
2.1	(Continued) Details of the papers reviewed.	20
2.1	(Continued) Details of the papers reviewed.	21
2.2	Studies contributing in the data gathering step. Studies gathering data with a sub-daily resolution are considered as <i>high-resolution</i> , <i>low-resolution</i> otherwise.	25
2.3	Studies contributing in the water end-uses characterization step.	29
2.4	Studies contributing in the user modeling step. Legend for multivariate analysis approaches: E = economic-driven; GS = geo-spatial; P = psychographic driven; AR = autoregressive. Legend for behavioural models approach: single = single user model; multi = multi-user model.	34
2.4	(Continued) Studies contributing in the user modeling step.	35
2.5	Studies contributing in the personalized Water Demand Management Strategies (WDMSs) step. Different WDMSs are considered: E = educational; F = financial; L = legislative; M = maintenance; T = technological.	39
2.6	Main research challenges for the use of smart meters in residential water demand modeling and management.	42
3.1	Fraction of energy assigned to each appliance (\hat{h}_i) and actual fraction of energy consumed by each appliance (h_i). Results obtained by using the optimization-based algorithm presented in the paper and the FHMM-based approach.	69
3.2	Relative Square Errors and R^2 coefficients. Results obtained by using the optimization-based algorithm presented in the paper and the FHMM-based approach.	69
3.3	Disaggregation performance metrics for HSID tested on a house comprising 11 appliances from the REDD dataset. Fields filled with "-" represent R^2 values lower than 0.	81

List of Tables

3.4	F-score for supervised HSID applied to power load data with different data sampling resolutions.	84
3.5	PCE (%) for supervised HSID applied to power load data with different data sampling resolutions.	84
3.6	R2 for supervised HSID applied to power load data with different data sampling resolutions. Fields filled with "-" represent R2 values lower than 0.	84
3.7	F-score for supervised HSID applied to water data with different data sampling resolutions.	92
3.8	PCE (%) for supervised HSID applied to water data with different data sampling resolutions.	92
3.9	R2 for supervised HSID applied to water data with different data sampling resolutions. Fields filled with "-" represent R2 values lower than 0.	93
4.1	Water consumers' and household features list.	111
4.2	Naive Bayes Classifier prediction accuracy based on Feature Selection (FS) algorithms outputs.	118
4.3	J48 Decision Tree prediction accuracy based on FS algorithms outputs.	118
5.1	Characterization of main heterogeneous typologies of behaviors.	136
5.2	Factor mapping results for customers' clusters to target for water demand management programs.	141

Acronyms

A

AEFI Actual Energy Fraction Index. 66, 69
ANN Artificial Neural Network. 49

B

BLOGREG Bayesian Logistic Regression. 108

C

CFS Correlation Feature Selection. 108

D

DSM Demand-Side Management. 10, 14, 123–125, 128, 135, 137, 139, 141
DTW Dynamic Time Warping. 28, 49, 50, 60–62

E

EEFI Estimated Energy Fraction Index. 65–67

F

FCBF Fast Correlation Based Filter. 108
FHMM Factorial Hidden Markov Model. 13, 44, 47, 50, 57, 59–63, 67–70, 72–75, 77, 80, 83
FS Feature Selection. 107, 118
FW Feature Weighting. 107, 119

H

HMM Hidden Markov Model. 28, 49, 59, 67
HSID Hybrid Signature-based Iterative Disaggregation. 13, 44, 47, 50, 51, 56–58, 62, 63, 70–73, 75–87, 89, 91–95

I

ISDTW Iterative Subsequence Dynamic Time Warping. 13, 44, 47, 50, 62, 63, 70, 73

N

NILM Non-Intrusive Load Monitoring. 13, 30, 43, 45–47, 49–51, 59, 71, 72, 95, 96

NILMTK Non-Intrusive Load Monitoring Toolkit. 67

P

PCA Principal Component Analysis. 101, 103, 126, 128

PCE Power Contribution Error. 71–73, 75, 77–82, 84, 92

PRIM Patient Rule Induction Method. 129

R

RSE Relative Square Error. 66, 67

S

SBMLR Sparse Bayesian Multinomial Logistic Regression. 108

SDTW Subsequence Dynamic Time Warping. 60, 61, 70

SOD Sparse Optimization Disaggregation. 46, 49, 51, 63, 64, 67, 69, 73, 77, 86, 95

SU Symmetrical Uncertainty. 108

T

t-SNE t-Distributed Stochastic Neighbour Embedding. 128

TRSE Total Relative Square Error. 65

W

WDMS Water Demand Management Strategy. 4, 5, 7, 8, 10, 13, 18, 26, 31–33, 38, 39, 41, 97, 99, 117, 120, 143

1

Introduction

1.1 Residential water demand management

World's urban population is expected to raise from current 54% to 66% in 2050 and to further increase as a consequence of the unlikely stabilization of human population by the end of the century (Gerland et al., 2014). By 2030 the number of mega-cities, namely cities with more than 10 million inhabitants, will grow over 40 (UNDESA, 2010). This will boost residential water demand (Cosgrove and Cosgrove, 2012). Despite the largest share of water generally goes to the agricultural sector, residential water demand currently has already an important role. It nowadays covers a large portion of the public drinking water supply worldwide (e.g., 60-80% in Europe (Collins et al., 2009), 58% in the United States (Kenny et al., 2009)), and is going to achieve more and more relevance with the rising of water demand, especially in densely populated urban areas.

This future scenario is going to pose several environmental and operational challenges. The concentration of the water demands of thousands or millions of people into small areas will considerably raise the stress on finite supplies of available freshwater (McDonald et al., 2011a). Besides, climate and land use change will further increase the number of people facing water shortage (McDonald et al., 2011b). Moreover, the expansion of urban areas, as well as an increase in urban population densities, is going to modify water demand in terms of magnitude, peak intensity, spatial and temporal distribution, and share between indoor and outdoor usage, requiring an adaptation of traditional supply and management schemes to meet future demands and prevent unsustainable resources exploitation.

In such a context, water supply expansion through the construction of new

storage and distribution infrastructures might be an option to escape water stress in some situations. Yet, geographical or financial limitations, as well as social, environmental and economic impacts of infrastructural interventions, largely restrict such options in most countries (McDonald et al., 2014). For instance, Gleick et al. (2003) claimed that approximately one-third of California's urban water use at that time could cost-effectively be saved without requiring new sources of water supply or big infrastructure interventions.

Here, acting on the demand side of water management represents a key alternative strategy for securing reliable water supply and reducing water utilities' costs (Beal et al., 2016). The implementation of several water saving programs, especially in the last two decades, demonstrated how WDMSs can complement supply-side management in order to meet future demands, and pursue water conservation levels, for instance through the definition of water conservation targets to save water in areas affected by prolonged periods of resource scarcity such as California (Arbués et al., 2003; Cahill and Lund, 2012) or recurrent low-recharge periods, like in Australia (Beal et al., 2016). Beside top-down water use limitation and conservation targets imposed by local governments and institutions, quick technology developments have fostered the actuation of efficiency interventions, consisting in the retrofitting of old water-consuming devices and replacement with higher efficient ones (e.g., low-flush toilets, efficient showerheads, efficient clotheswasher and dishwasher). Worldwide experiences have been proving that these interventions potentially lead to significant reductions in residential water consumption (Mayer et al., 2003, 2004; Lee et al., 2011). For instance, DeOreo et al. (2016) recently analyzed the indoor water consumption of a sample of single-family homes in the US, noticing a decrease of 22% in the average per-household water consumption, mainly because of improved efficiencies of clothes washers and toilets.

However, the extent of WDMSs' scope and objectives is not merely limited to water conservation. Demand side management interventions constitute also a cost-effective tool to support utilities' business (Gleick et al., 2003) and reduce their operational short and long term costs, e.g., through peak consumption reduction or time shift (Cardell-Oliver et al., 2016; Beal et al., 2016), motivated by rising infrastructure and operating costs to supply water to residential users (Commission et al., 2013; Cardell-Oliver et al., 2016). Finally, WDMSs can foster water distribution and consumption efficiency, meaning that resources waste for supply and use services are reduced (e.g., through improved leakage detection and maintenance of the water distribution network), while the quality of these services is secured.

These multi-fold opportunities offered by WDMSs are an intrinsic consequence of both the number of stakeholders involved in the design of such interventions (i.e., private and public water utilities, municipalities, national and

local governments and authorities), and the heterogeneity of those actors, (i.e., residential water consumers), whose demand reduction/change is the target of WDMSs. Householders, utilities, and public authorities do have different, potentially conflicting, objectives. For instance, householders can be interested in satisfying their residential water needs and pursue happiness, while reducing costs; utilities can be interested in minimizing their costs and optimizing network operation; public authorities do have to safeguard the quality of the service, while keeping into account environmental regulation and overall benefits for the society. Therefore, different combinations of WDMSs are likely to differently balance these objectives and heterogeneously distribute benefits among such actors. Additionally, we can think of water consumers as agents who decide when and how much water to use during the day, as well as for which specific end-use, depending on their needs. Consequently, the achievement of WDMSs objectives depends on the aggregate effect of actions undertaken by several independent and diverse agents. This picture suggests that a major role contributing to residential water demand is played by the behavioral component of water consumers. Thus, given the potentially different water consuming behaviors and consumers' socio-demographic characteristics within a community of users, a portfolio of customized and flexible water demand management interventions, differentiated among diverse groups of users, can lead to a better achievement of management objectives. The portfolio of customized WDMSs include those involving educational aspects, financial schemes aimed at incentivizing/discouraging specific behaviors, e.g., time-of-use tariffs (Cole et al., 2012), and attitude changes driven by water awareness. Indeed, this has been recognized also by international organizations: the goals set in the 2030 Agenda for Sustainable Development (Assembly, 2015) necessitate not only technological progress, but also changes in the attitude of citizens towards our shared natural resources. As a proof, the number of initiatives (e.g., UN World Water Day) organized every year by numerous Public Organizations, demonstrates the need of making citizens more aware of the impacts of their daily behavior and of stimulating long-term behavioral change towards a more sustainable use of natural resources.

1.2 Residential water-energy nexus

Many recent studies explored the value of considering mutual interactions between water and energy to design coordinated planning and management actions. In summary, this is motivated by the fact that water and energy are interconnected, in so far that water needs energy and energy needs water (Escrivà Bou, 2016). Energy production needs water: IEA (2012) estimates that energy production consumes 15% of water withdrawals. This because many

energy production processes either require water directly for energy production (e.g., hydropower production) or indirectly for secondary processed (e.g., thermo electric cooling processes). In turn, energy is needed along the whole water cycle: extraction and pumping, treatment, distribution, heating (alternatively to gas), use, and recycling. Also, energy is needed for the maintenance of water distribution networks. In California, for instance, the California Energy Commission 2005 estimated that approximately 20% of the state's annual electricity production and 30% of the state's annual natural gas production are used as water-related energy in the water cycle.

Considering the water-energy nexus in the residential sector, and more specifically the interactions between water and energy at the single- household level (thus excluding the the amount of energy required for water extraction, pumping and delivery) most of the nexus has to be attributed to water heating. Ryan et al. (2010) state that about 17% of a household energy use is due to water heating (Abdallah and Rosenberg, 2014). Just to provide the reader with a more practical example, the US Environmental Protection Agency states that heating by electricity the water needed for a 5-minute faucet use requires approximately the same amount of energy used in 14 hours by a 60-watt bulb.¹ While traditionally management practices have been designed separately for residential water and energy demand and supply, only recent studies (e.g., Abdallah and Rosenberg, 2014) explored the linkages between household water and energy demands, as well as their implication and benefits for the design of coordinated management actions.

With the development of this thesis, we would like to stress two additional points related to the water-energy nexus in the residential sector, and in particular regarding the opportunities it offers to the development of users' behavioral models. First, beside the direct link between water and energy hidden in the water-related energy use, residential water and energy demands are indirectly related through consumers' behavioural factors. As mentioned above, household water consumption depends on a number of socio-psychographic variables. This stands also for energy demand. Moreover, typical household water and energy consumption patterns usually reflect users' habits and routines (if any) and reveal users' attitudes and preferences towards consumption. Indeed, to our knowledge, no studies in the literature have explored these indirect behavioural linkages between water and energy consumption habits and drivers. Identifying residential water and electricity consumers' routines and understanding weather the typical consumption patterns of water and energy consumption present both similarities (e.g., synchronized demand peaks) and synergies (e.g., water-related energy use) that depend on users' behaviors, characteristics, and attitudes can foster the identification of groups of users to target for

¹<https://www3.epa.gov/region9/waterinfrastructure/waterenergy.html>

1.3. Opportunities and challenges for modelling residential water (and energy) demand

joint water and energy conservation, behavioural change, or more widely management actions. Second, smart electricity grids (Kramers et al., 2014; Niese et al., 2014) started developing prior than the first deployments of smart water meters. From a methodological point of view, this is an opportunity for the development of residential water demand models. Indeed, several challenges of modelling consumers' behavior, for instance for the study of single-appliance consumption contribution or for the extraction of typical load shapes, have been largely studied in the energy sector, pushing towards the development of algorithms and models able to manage smart metered data. The adaptation of promising methods to water applications, as well as the development of new modelling methodologies able to deal both with residential water and energy, possibly considering synergies between the two, represent cost-effective paths to advance research in both directions, to provide supporting tools to multi-utilities, and to support the design of coordinated management strategies.

1.3 Opportunities and challenges for modelling residential water (and energy) demand

In recent years a number of WDMSs has been developed and applied worldwide (for a review, see Inman and Jeffrey, 2006; Cominola et al., 2015a, and references therein). However, the effectiveness of WDMSs is often context-specific and strongly depends on our understanding of the drivers inducing people to consume or save water (Jorgensen et al., 2009). Therefore, models that quantitatively describe how water demand is influenced and varies in relation to exogenous uncontrolled drivers (e.g., seasonality, climatic conditions), water consumers' and household characteristics (e.g., number of house occupants, demographic features, house size), and demand management actions (e.g., water restrictions, pricing schemes, education campaigns) are essential to explore water users' response to alternative WDMSs, ultimately supporting strategic planning and policy design.

Traditionally, water demand models focus on different temporal and spatial scales. At the lowest resolution, studies have been carried out, mostly in the 1990s, to model water demand at the urban or block group scale, using low time resolution (i.e., above daily) consumption data retrieved through billing databases or experimental measurement campaigns on a quarterly or monthly basis. The main goal of these works is to inform regional water systems planning and management on the basis of estimated relationships between water consumption patterns and socio-economic or climatic drivers (e.g., House-Peters and Chang, 2011). Such works improved our understanding of urban water use, especially because they first pointed out that urban water use is the outcome of a combination of heterogeneous drivers and psychographic variables, namely

demography, behaviors, technology, structural and environmental factors. Yet, data resolution and availability constitute main limitations for the wide usability of these models towards the customization of WDMSs. On the one hand, low spatial and temporal resolution of water consumption data, as metered with traditional water meters, does not allow a detailed characterization of individual user's daily consumption patterns or an accurate breakdown of such uses among different indoor and outdoor end-uses. On the other hand, also the number and spatial resolution of psychographic variables, usually limited to a few statistics describing groups of users larger than individual households, entailed *a priori* assumptions on the candidate drivers to consider for model building and limited their suitability to represent individual water consumers' heterogeneity.

The advent of smart meters (Mayer and DeOreo, 1999) in the late 1990s made available new water consumption data at very high spatial (household) and temporal (from several minutes up to few seconds) resolution, enabling the development and application of data analytics tools and mathematical models to accurately characterize sub-daily water consumption patterns, as well as end-use water consumption profiles. Similarly to the recent developments in integrated smart solutions (Laniak et al., 2013; Hilty et al., 2014), the use of smart meters provides essential information to construct models of the individual consumers behaviors, which can be employed for designing and evaluating consumer-tailored WDMSs that can more effectively modify the users' attitude favoring water saving behaviors. Indeed, the ability to monitor water consumption at the single-household scale and with a sub-daily sampling resolution pushed the development of behavioural change programs promoting water saving attitudes via tailored feedbacks (Fielding et al., 2013). Beside the fact that smart water meters themselves constitute technologies that promote water conservation awareness, a relevant water saving potential can be obtained providing feedback to users about their water consumption or suggestions on customized water savings practices, especially through electronic visual and alarming monitoring devices (e.g., Willis et al., 2010; Froehlich et al., 2012; Desley et al., 2013; Sonderlund et al., 2014). Benchmark experiences in the water sector show that behavioral change interventions can lead to reductions of overall water consumption between 3 and 10% (Inman and Jeffrey, 2006; Willis et al., 2010; Anda et al., 2013; Davies et al., 2014; WaterSmart software, 2014), with some exceptional reductions between 15 and 27% (Mayer et al., 2004; Willis et al., 2010), the latter mainly due to interventions on specific end-uses joined with an increased efficiency of water consuming devices. This appears also in line with the literature on energy savings, as customized feedbacks have been shown to allow energy usage reductions between 2 and 20% (Darby, 2006; Fischer, 2008; Ehrhardt-Martinez et al., 2010).

Regarding the identification of exogenous and individual user's drivers of

residential water consumption, as well as the characterization of water users through a number of socio-psychographic variables, recent studies (e.g., Makki et al., 2013) advanced correlation analysis to determine the influence of individual household factors, such as demographic (e.g., number of house occupants, age), house (house size and age), economic (e.g., income), and personal (education, attitudes) characteristics on water consumption. Yet, major data gathering limitations still exist, as socio-psychographic data are generally collected with *ad hoc* surveys and interviews to individual users, making the whole process time and resources demanding and hard to replicate with high frequency. However, the recent development of online web platforms, such as the SmartH2O user portal², which allow users to create their own household profile, inserting interactively information on their household, demography, water consuming fixtures, etc. represents a promising step forward for the deployment of users' data gathering systems more accessible both to users themselves and utilities, easy to update, and less intrusive (users can optionally insert their information whenever they want to and do not have to attend interviews in person).

Overall, water consumers models benefit from the quick technological and computational developments allowing for metering and gathering water consumers' data with higher resolutions, sampling frequencies, and quicker update time. Indeed, such progress is making available an amount of data that was not available just a few years ago, and is going to increase in the next future with the large scale deployment of smart metering systems and the availability of data from new data sources. Yet, apart from the technical challenges posed by the storage and management of big datasets, a paramount challenge present in all the fields where big data and big data analysis are involved is how to turn these data into useful information.³ In the urban water research field this means how to extract from big smart metered datasets, represent, and visualize relevant information supporting the design of customized water demand management strategies. This motivation calls urban water researchers and modellers to the development of data mining techniques able to perform the above mentioned tasks related to information extraction in a fast, automatic, reproducible and scalable way.

1.4 Thesis motivation, objectives, and scope

Given the above mentioned research opportunities and challenges for modelling residential water demand, in this thesis, we contribute novel methodologies and data mining tools for profiling, analyzing, and modeling residential

²SmartH2O official website: www.smarth2o-fp7.eu.

³On this topic, see a recent post by Julia Wagemann, Technical Data Analyst at the European Centre for Medium-Range Weather Forecasts <http://blogs.ecu.eu/divisions/essi/2016/07/11/big-earth-science-data-boon-or-bane/>.

water consumers based on high temporal and spatial resolution data of residential water consumption, collected by smart meters, and also coupled with several qualitative and quantitative data describing consumers' psychographic features. The main focus is on water demand modelling and management. However, in this thesis we either assess the inter-portability of the methodologies between the two fields of applications (i.e., water and energy), or present integrated water-energy applications. This research is motivated by the opportunities mentioned in the previous sections of this introduction and, above all, by the need of mathematical tools that are:

- suitable to manage big smart metered water and energy datasets, comprising several hundreds or thousands of residential users;
- able to automatically extract relevant information on users behavior out of these data, possibly in conjunction with users' data from other sources (e.g., surveys, social computing, online platforms);
- flexible enough to be usable, portable, or easily adaptable both for water and energy applications, as well as integrated water-energy applications;
- able to support and inform the design of customized water and energy Demand-Side Management (DSM) interventions.

Driven by the above mentioned challenges, the main goal of this thesis is advancing data analysis and mathematical models to extract information on water consumers' behavior out of smart-metered data, ultimately informing and proposing recommendations to customized demand management. Generally, residential water demand modelling and management studies can be organized in a simple four-steps framework, i.e., (i) data gathering, (ii) water end-uses characterization, (iii) user modeling, and (iv) design and implementation of personalized WDMSs, as reported in Figure 1.1. We provide a more detailed description of each of these steps in Chapter 2, where we identify a common framework for the critical review of the best practices currently existing in the research field of residential water demand modelling and management. More specifically, this thesis aims at contributing in different phases of the identified framework, i.e., 2 to 4 in Figure 1.1, while for the data gathering phase we rely on existing smart metering technology. Our first goal is advancing end-use characterization studies, assessing the value of data sampling resolution and similarities/differences between water and energy data to foster cost-effective and accurate methodologies for end-use disaggregation. Second objective is searching for water consumers' habits and routines, as well as identifying the motivations, i.e., the most relevant determinants, of their water consuming or saving behaviors, by mining smart-metered data gathered at the single-household level. Finally, our ultimate goal is exploring how the outcomes of smart-metered data

analysis and mathematical models can inform the design of water and energy demand management strategies, as well as the identification of groups of users to be targeted for such actions, and customized recommendations.

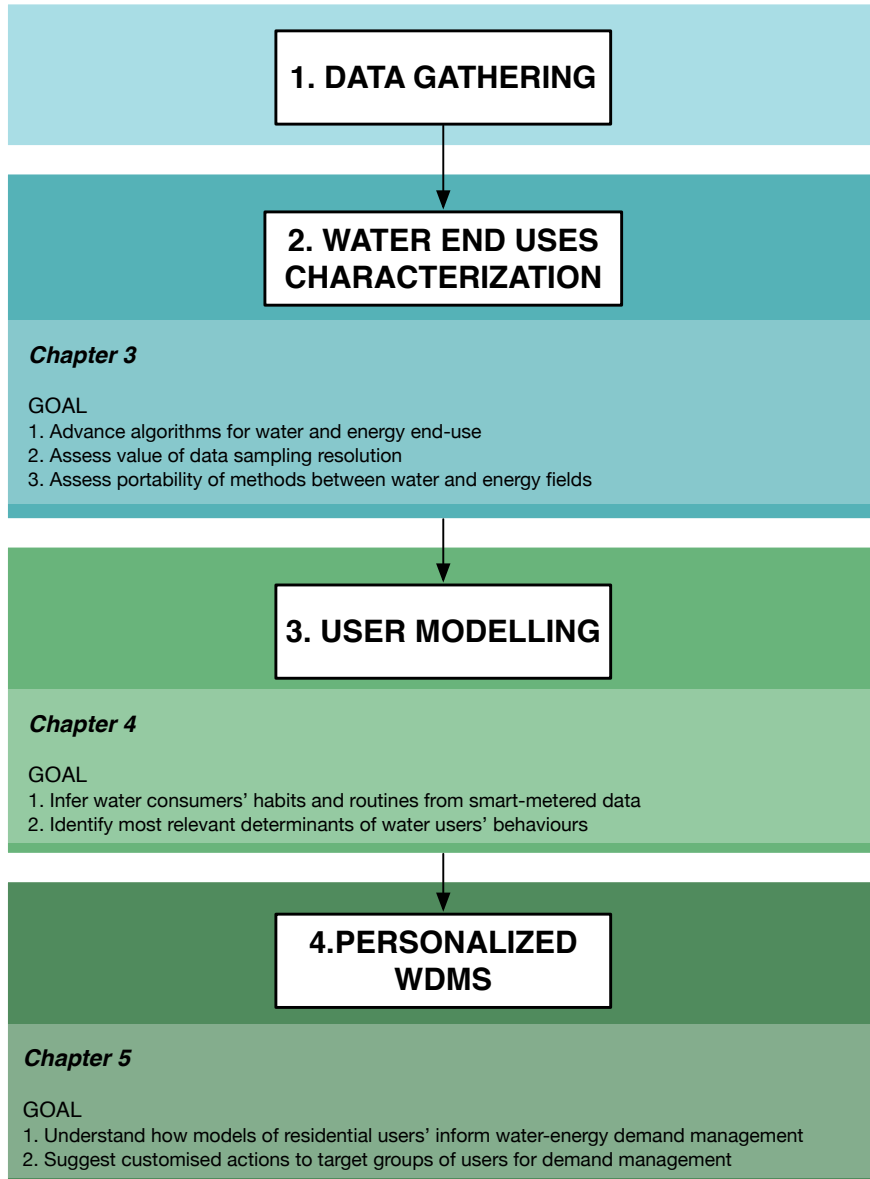


Figure 1.1: Simplified framework reporting the phases of residential water demand modelling and management studies. The objectives of core chapters are briefly reported in each box.

Overall, the high-level objective of this research is not limited to providing modelling tools that contribute to each phase of those in Figure 1.1 in accordance to the mentioned specific objectives. Most of all, our target is to integrate such tools within a common methodological framework and, thus, deliver a suite of data analysis and modelling tools supporting the decisional processes needed for the design of customized demand-side management strategies. We thus carried out this research within the following scope:

- this research considers only water and electricity demand data monitored in the residential sector. Demands from other sectors, e.g., commercial, does not fall within the scope of this research;
- we developed most of this research as part of the SmartH2O project⁴, a European project funded by the European Commission, and member of the ICT4Water cluster⁵. The goal of SmartH2O is developing an ICT platform for improving the management of urban and peri-urban water demand thanks to the integrated use of smart meters, social computation, and dynamic water pricing, based on advanced models of consumer behavior;
- despite the support of the SmartH2O project, this study is not limited to the contexts of the SmartH2O case studies, but several case studies worldwide, including Australia and the United States are included;
- given the presence of multiple case studies comprising data with different sampling resolution, the scope of this study is not limited to the analyses of data gathered with a specific sampling frequency. The analyses and methods proposed in this research are adapted in order to be able to accommodate data with diverse resolutions, depending on their availability.

The innovative aspects of this research are two-fold. First, innovation is present in the methodologies we present, partially developed by adapting, further extending and integrating data mining techniques from several data mining and modelling research fields, i.e., signal processing, information extraction, data reduction, and customer behaviour studies. Secondly, the set of modelling tools we develop in this research cover all the phases of residential water demand modeling and management, from big smart metered data processing, to end-uses classification, user modelling and customer segmentation and recommendations to water (and electricity) demand management. Such integration can constitute a decision support system for urban water and electricity demand management and planning. Further details about the research method and the content of each chapter are reported in the next section.

⁴<http://www.smarth2o-fp7.eu/>

⁵<http://www.ict4water.eu/>

1.4. Thesis motivation, objectives, and scope

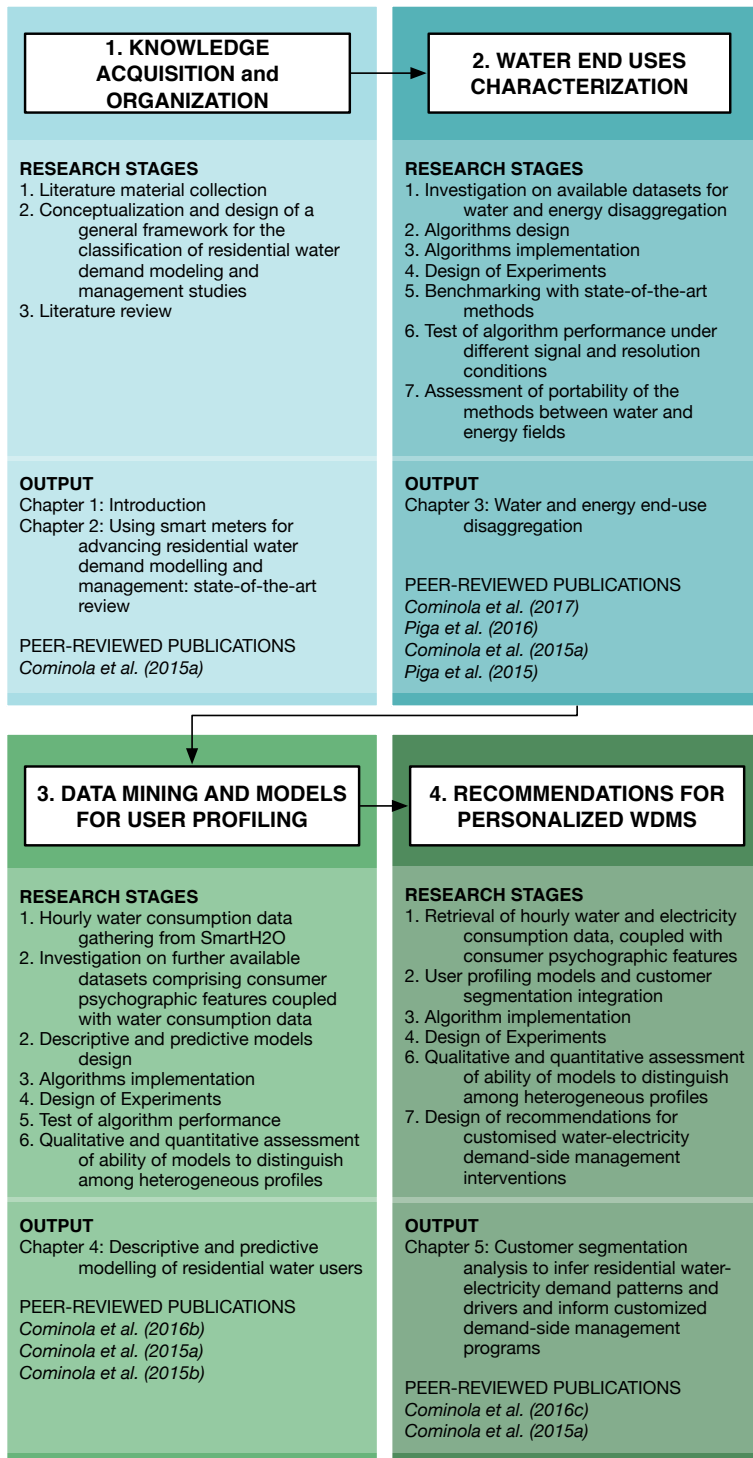


Figure 1.2: Overview of the research method adopted in this thesis.

1.5 Research Method

The research developed along this multi-disciplinary project followed the research method we represent in summary in Figure 1.2. It is composed by four main phases, i.e., (i) knowledge acquisition and organization, (ii) water end uses characterization, (iii) data mining and models for user profiling, and (iv) recommendations for personalized WDMS. In the figure, each phase of the research method we adopted is characterized in terms of activities, i.e., methodological research stages, and output, i.e., peer-reviewed publications and written contributions to this thesis.

As stated in the previous section with the objectives of this work, we aimed at delivering an integrated suite comprising diverse modelling tools able to inform and support each phase of the framework identified for residential water demand modelling and management studies and project (Figure 1.1). In the attempt to do so, our research approach is based on two pillars, i.e., (i) a detailed and exhaustive knowledge acquisition from the state-of-the-art literature regarding each phase of the framework and (ii) the contribution of tools and publications for each phase of the framework, later integrating such tools as the final deliverable of this research. Following the first pillar, we were able to publish a comprehensive literature review (see Cominola et al. (2015a)) which comprises the critical analysis of works related to each of the methodological steps. For this reason, the reference is reported as an output for all the four boxes in Figure 1.1. In addition, we exploited the knowledge acquired along with this first phase to better define our research objectives, identify the relevant challenges commented in the previous section, and set the requirements for our methodological contributions. Regarding the methodological contributions of this thesis, each phase of our research method included several sub-tasks, spanning from data gathering or retrieval, to the quantitative and qualitative assessment of the models we developed. On the attempt of pursuing our ultimate target and integrate all the methods we designed and implemented, a major challenge has been often represented by data availability and accessibility. The data we had access to from the SmartH2O project or other external sources presented some limitations for an integrated application of all the methods developed on a common dataset. Indeed, for this reason we report a data gathering/retrieval stage for each of the phases in Figure 1.2 and, for instance in the case of end-use disaggregation, we first developed model prototypes based on electricity data and then adapted them to be suitable also for water applications.

Finally, regarding the organization of the outputs listed in the figure in terms of written contribution in this thesis, after this introduction, each chapter relates to a specific methodological phase, presenting tools and numerical results. More details on the organization of each chapter are reported in the next section.

1.6 Thesis outline

1.6.1 Chapter 2

In Chapter 2, we propose both a methodological contribution, as well as a review of the state-of-the-art literature on residential water demand modeling and management. Firstly, we set up a general framework for the classification of residential water demand modeling and management studies (see Figure 2.2). According to each methodological phase of the proposed framework, i.e., (i) data gathering, (ii) water end-uses characterization, (iii) user modeling, and (iv) design and implementation of personalized WDMSs, we then comprehensively review and classify 134 studies that in the last 25 years advanced methodologies and tools in the residential water research domain. In particular, our review analyzes consolidated approaches, describes emerging trends, and identifies potential future developments in this continuously and quickly evolving field.

The content of this chapter is adapted from Cominola et al. (2015a).

1.6.2 Chapter 3

In Chapter 3 we deal with the problem of disaggregating the total household water (or power) consumption, metered by a single-point meter with sub-daily resolution, into single-appliance consumption patterns. We contribute two algorithms for water and energy end-use disaggregation. The first algorithm addresses the disaggregation problem as a least-square error minimization problem, with an additional (convex) penalty term aiming at enforcing the disaggregate signals to be piecewise constant over the time. The second algorithm, called HSID, is based on the combination of FHMMs, which provides an initial approximation of the end-use trajectories, and Iterative Subsequence Dynamic Time Warping (ISDTW), which processes the end-use trajectories in order to match the typical power/water consumption pattern of each appliance. As the literature on Non-Intrusive Load Monitoring (NILM) techniques for power load disaggregation is more consolidated than the more recent literature on water end-use disaggregation, we first test the two algorithms onto real-world power load data and against benchmark methods from the literature and, secondly, we extend the application of HSID to the disaggregation of synthetic water data at multiple sampling resolutions. HSID is also demonstrated to be robust with respect to noisy signals, scalable to dataset including a large set of appliances and can be successfully used in non-intrusive experiments without requiring appliance-level measurements.

Some sections of this chapter are adapted from Piga et al. (2016), Cominola et al. (2017) and Cominola et al. (2016a).

1.6.3 Chapter 4

In Chapter 4 we contribute two novel modelling tools to infer water consumers' habits and routines, as well as identify the most relevant determinants of their water consuming or saving behaviors, at the household level. The use of such models is key to characterize the heterogeneity of water users' behaviors in a community, predict urban water demand variability in space and time, and explore the effects of different WDMSs for the residential sector. More specifically, we first develop a descriptive modeling procedure, based on a combination of clustering and principal component analysis, which allows performing water users' segmentation on the basis of their eigenbehaviors (i.e., routines, recurrent water consumption behaviors) automatically identified from smart metered consumption data. We test our approach onto a dataset of water consumption data metered with hourly resolution from 175 households in the Swiss Municipality, retrieved as part of the SmartH2O project, demonstrating the suitability of the method for identifying typical profiles of water consumption, as well as for clustering water users' in distinct groups. Secondly, we propose a predictive modeling approach, consisting of a two-step procedure that (i) extracts the most relevant determinants of users' consumption and (ii) identifies and learns a predictive model of water consumers' profile. We test the method considering a real case study with low resolution data of over 1500 water users from the Kimberley and Pilbara Western Australia. Results show the effectiveness of the proposed method in capturing the influence of candidate determinants on residential water consumption, as well as in attaining sufficiently accurate predictions of users' consumption profiles.

The content of this chapter is adapted from Cominola et al. (2016b) and Cominola et al. (2015b).

1.6.4 Chapter 5

In Chapter 5 we integrate part of the components of the models described in Chapter 4 and contribute a customer segmentation analysis of over 1000 residential water and electricity accounts in South California through data mining techniques, with the goals of exploring differences between typical residential water and electricity demand patterns, interpreting these typical demand patterns in terms of users' behavior, and providing useful insights to coordinated water-energy demand management programs based on resources conservation or demand peak shifting. This contribution is based on (i) the fact that demand-management benefits from tools that discriminate among heterogeneous daily water/energy consumption profiles, as well as (ii) the water-energy nexus in the residential sector, whose implications towards resources conservation and efficient use are pushing towards synergies for the development of coordinated

water-energy demand management actions to meet future water and energy demands, and reduce utilities' costs. Our procedure combines eigenbehavior extraction for the identification of typical consumers profiles from smart metered hourly water-energy time series, clustering for customer segmentation based on profile similarities, and factor mapping to infer the potential determinants of targeted profiles to inform the design of customized DSM interventions. Our procedure captures heterogeneous water-electricity consumption profiles of the considered accounts, highlighting differences between daily time-of-use of water and electricity, and allowing for the characterization of accounts based on psychographic and behavioral factors. Finally, we design recommendations for water-electricity programs targeting resources conservation or demand peak shifting.

The content of this chapter is adapted from (Cominola et al., 2016c), in preparation.

1.6.5 Chapter 6

In Chapter 6 we summarize the achievements of this PhD thesis, gaps and ideas raised for future research, and general conclusions.

2

Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

Abstract¹

Over the last two decades, water smart metering programs have been launched in a number of medium to large cities worldwide to nearly continuously monitor water consumption at the single household level. The availability of data at such very high spatial and temporal resolution advanced the ability in characterizing, modeling, and, ultimately, designing user-oriented residential water demand management strategies. Research to date has been focusing on one or more of these aspects but with limited integration between the specialized methodologies developed so far. In this chapter, we present a comprehensive review of the literature in this quickly evolving water research domain. The review contributes a general framework for the classification of residential water demand modeling studies, which allows revising consolidated approaches, describing emerging trends, and identifying potential future developments. In

¹The content of this chapter has been adapted from the following paper: Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2015a. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environmental Modelling & Software* 72, 198–214.

particular, the future challenges posed by growing population demands, constrained sources of water supply and climate change impacts are expected to require more and more integrated procedures for effectively supporting residential water demand modeling and management in several countries across the world.

2.1 Introduction

As mentioned in Chapter 1, a variety of WDMSs has been proposed and applied in recent years (for a review, see Inman and Jeffrey, 2006, and references therein), in order to complement supply-side management interventions to secure reliable water supply and reduce water utilities' costs (Gleick et al., 2003). However, the effectiveness of these WDMS is often context-specific and strongly depends on our understanding of the drivers inducing people to consume or save water (Jorgensen et al., 2009). Models that quantitatively describe how water demand is influenced and varies in relation to exogenous uncontrolled drivers (e.g., seasonality, climatic conditions) and demand management actions (e.g., water restrictions, pricing schemes, education campaigns) are essential to explore water users' response to alternative WDMS, ultimately supporting strategic planning and policy design.

A general procedure to study residential water demand management relying on the high-resolution data nowadays available can be structured in the following four phases (see Figure 2.2): (i) data gathering, (ii) water end-uses characterization, (iii) user modeling, (iv) design and implementation of personalized WDMSs. In the literature, a number of tools and techniques have been proposed for each of these steps, with many works focused either on the data gathering process (e.g., Cordell et al., 2003; Boyle et al., 2013) or on the analysis of WDMSs (e.g., Inman and Jeffrey, 2006). Yet, to our knowledge, a systematic and comprehensive review of residential water demand modeling and management is still missing. This review contributes the first effort of classification and critical analysis of 134 studies that in the last 25 years (Figure 2.1) contributed new methodologies and tools in one or more of the steps of the above procedure (see Table 2.1).

This chapter is structured according to the procedure shown in Figure 2.2: the current status, research challenges, and future directions associated to each phase are discussed in Sections 2.2-2.5, while the last section reports final remarks and directions for follow up research.




Figure 2.1: Five-years count of the 134 publications reviewed in this study.

Table 2.1: Details of the papers reviewed.

Reference	Location	Data gathering	Water end-uses	User modeling	Personalized WDMS
Anda et al. (2013)	Australia	x			
Boyle et al. (2013)	N/A	x			
Willis et al. (2013)	Australia	x		x	
Froehlich et al. (2011)	N/A	x	x		
Wong et al. (2010)	Hong Kong	x			
Froehlich et al. (2009)	N/A	x			
Kim et al. (2008)	N/A	x			
Heinrich (2007)	New Zealand	x	x		
Olmstead et al. (2007)	USA	x		x	
Kowalski and Marshallsay (2005)	UK	x	x		
Evans et al. (2004)	N/A	x			
Mayer et al. (2004)	USA	x	x		x
Mori et al. (2004)	N/A	x			
Cordell et al. (2003)	Australia	x			
Sanderson and Yeung (2002)	N/A	x			
Mayer and DeOreo (1999)	USA	x			x
Nguyen et al. (2014)	Australia		x		
Nguyen et al. (2013a)	Australia		x		
Nguyen et al. (2013b)	Australia		x		
Cardell-Oliver (2013a)	Australia		x		
Cardell-Oliver (2013b)	Australia		x		
Aquacraft Inc. (2011)	USA		x		
Beal et al. (2011a)	Australia		x		
DeOreo et al. (2011)	USA		x		
Mead and Aravinthan (2009)	Australia		x		
Willis et al. (2009a)	Australia		x		
Willis et al. (2009b)	Australia		x		
Roberts (2005)	Australia		x		x
Kowalski and Marshallsay (2003)	UK		x		
Loh et al. (2003)	Australia		x	x	
DeOreo and Mayer (2000)	USA		x		
DeOreo et al. (1996)	USA		x		
Mayer and DeOreo (1995)	USA		x		
DeOreo and Mayer (1994)	USA		x		
Makki et al. (2015)	Australia			x	
Beal et al. (2014)	Australia			x	

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

Table 2.1: (Continued) Details of the papers reviewed.

Reference	Location	Data gathering	Water end-uses	User modeling	Personalized WDMS
Kanta and Zechman (2014)	N/A			x	
Beal and Stewart (2014)	Australia			x	
Matos et al. (2014)	Portugal			x	
Talebpour et al. (2014)	Australia			x	
Romano et al. (2014)	Italy			x	
Cardell-Oliver and Peach (2013)	Australia			x	
Beal et al. (2013)	Australia			x	
Bennett et al. (2013)	Australia			x	
Cahill et al. (2013)	USA			x	
Cole and Stewart (2013)	Australia			x	
Makki et al. (2013)	Australia			x	
Beal et al. (2011b)	Australia			x	
Gato-Trinidad et al. (2011)	Australia			x	
Grafton et al. (2011)	10 OECD countries			x	
House-Peters and Chang (2011)	N/A			x	
Lee et al. (2011)	USA			x	
Nasseri et al. (2011)	Iran			x	
Qi and Chang (2011)	USA			x	
SDU (2011)	USA			x	
SJESD (2011)	USA			x	
Willis et al. (2011)	Australia			x	
Blokker et al. (2010)	Nederland			x	
Chang et al. (2010)	USA			x	
Lee and Wentz (2010)	USA			x	
Polebitski and Palmer (2010)	USA			x	
Rosenberg (2010)	Jordan			x	
Russell and Fielding (2010)	N/A			x	
Chu et al. (2009)	China			x	
Corbella and Pujol (2009)	N/A			x	
Fox et al. (2009)	UK			x	
Galán et al. (2009)	Spain			x	
Jorgensen et al. (2009)	N/A			x	
Olmstead and Stavins (2009)	N/A			x	
Praskievicz and Chang (2009)	Korea			x	
Balling et al. (2008)	USA			x	
Lee and Wentz (2008)	USA			x	
Alvisi et al. (2007)	Italy			x	
Balling and Gober (2007)	USA			x	
Gato et al. (2007)	Australia			x	
Rosenberg et al. (2007)	Jordan			x	
Wentz and Gober (2007)	USA			x	
Gato (2006)	Australia			x	
Altunkaynak et al. (2005)	Turkey			x	
Fullerton and Elias (2004)	USA			x	
Aly and Wanakule (2004)	USA			x	
Syme et al. (2004)	Australia			x	
Brookshire et al. (2002)	N/A			x	
Zhou et al. (2000)	Australia			x	
Zhou et al. (2002)	Australia			x	
Espey et al. (1997)	N/A			x	
Molino et al. (1996)	Italy			x	
Homwongs et al. (1994)	USA			x	
Lyman (1992)	USA			x	
Griffin and Chang (1991)	USA			x	
Rixon et al. (2007)	Australia			x	
Schneider and Whitlatch (1991)	USA			x	
Miaou (1990)	USA			x	
Maggioni (2015)	USA				x
Sonderlund et al. (2014)	N/A				x
Molinos-Senante (2014)	Spain			x	x
Britton et al. (2013)	Australia				x
Fielding et al. (2013)	Australia				x
Stewart et al. (2013)	Australia				x
Carragher et al. (2012)	Australia				x
Cole et al. (2012)	Australia				x
Froehlich et al. (2012)	USA				x
Froes Lima and Portillo Navas (2012)	Brazil				x
DeOreo (2011)	USA				x
Willis et al. (2010)	Australia				x
Mead and Aravinthan (2009)	Australia				x
Steg and Vlek (2009)	N/A				x
Britton et al. (2008)	Australia				x
Grafton and Ward (2008)	Australia				x
Worthington and Hoffman (2008)	N/A				x
Brennan et al. (2007)	Australia				x

Table 2.1: (Continued) Details of the papers reviewed.

Reference	Location	Data gathering	Water end-uses	User modeling	Personalized WDMS
Brooks (2006)	N/A				x
Hensher et al. (2006)	Australia				x
Inman and Jeffrey (2006)	N/A				x
Howarth and Butler (2004)	UK				x
Arbués et al. (2003)	N/A				x
Duke et al. (2002)	USA				x
Geller (2002)	N/A				x
Garcia and Thomas (2001)	France				x
Kanakoudis (2002)	Greece				x
Renwick and Green (2000)	USA				x
Renwick and Archibald (1998)	USA				x
Dandy et al. (1997)	Australia				x
Gurung et al. (2015)	Australia			x	x
Gurung et al. (2014)	Australia			x	
Suero et al. (2012)	USA			x	x
Giacomoni and Berglund (2015)	USA			x	x
Escriva-Bou et al. (2015a)	USA			x	x
Escriva-Bou et al. (2015b)	USA			x	x
Kenney et al. (2008)	USA			x	x
Kenney et al. (2004)	USA				x
Dalhuisen et al. (2003)	N/A			x	x
Mayer et al. (2003)	USA	x	x		x
Mayer et al. (2000)	USA	x	x		x

2.2 Data gathering

Residential water consumption data gathering (box 1 in Figure 2.2) represents the first step needed to build the baseline upon which the water demand is estimated and management strategies are designed. Depending on the sampling frequency, we distinguish two main classes, namely *low-resolution* and *high-resolution* data, which delimit the type of the analysis that can be performed.

2.2.1 Low resolution data

Periodically billed data are characterized by a low level of resolution and recording frequency. Although water consumption is detected with the precision of kilolitres, readings are generally recorded with the frequency of the quarter of year at most (Britton et al., 2008). This low resolution restricts the use of these data to regional planning, where statistical analysis estimating the amount of domestic water consumption can be used to forecast the aggregated water demand at the municipal or district level. In particular, such data have been widely used to study the effect of economic variables and seasonality on the water use at the regional scale since the seminal works by Howe and Linaweaver (1967); Young (1973); Berk et al. (1980); Howe (1982); Maidment and Parzen (1984); Thomas and Syme (1988) (for a review see House-Peters and Chang, 2011, and references therein). Those approaches relied on simple econometric models and time series models based on multivariate regression, and required limited datasets and low computational resources. Their main drawback is related to their limited capability of representing the spatial and temporal heterogene-

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

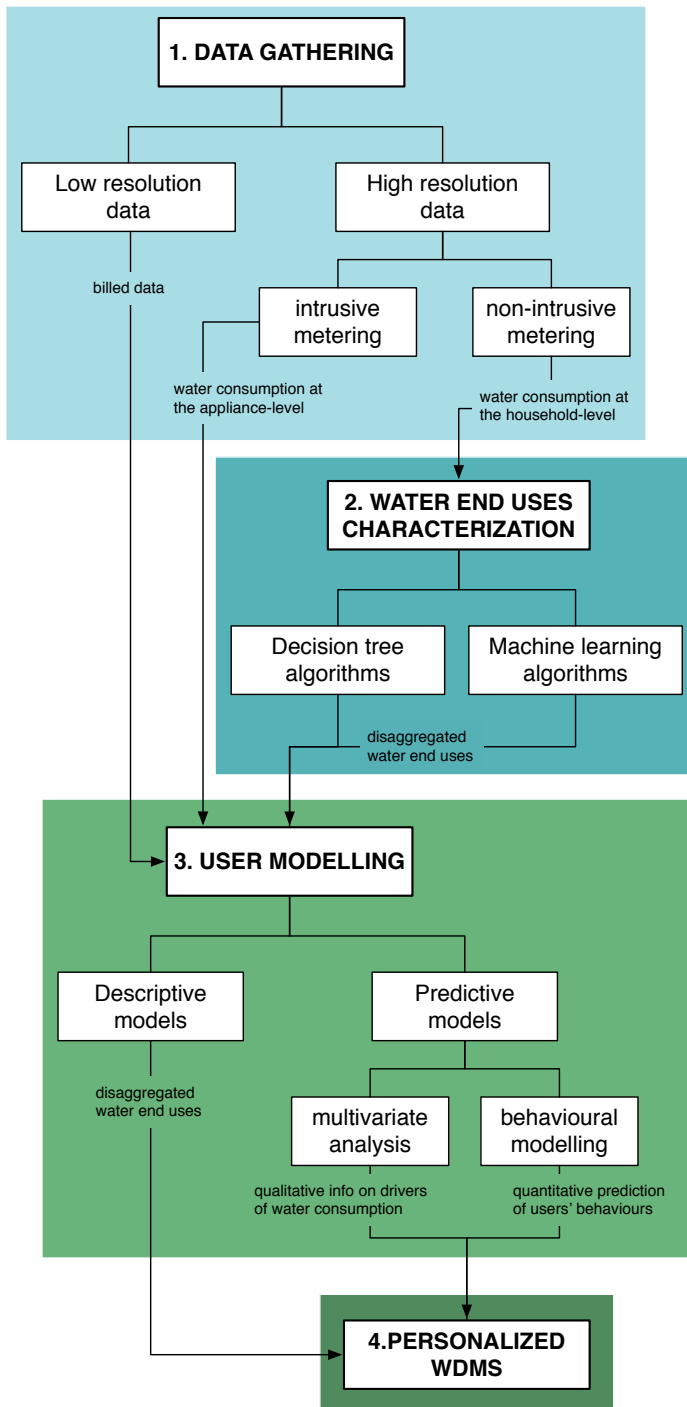


Figure 2.2: Flowchart of the general procedure for studying residential water demand management.

ity of residential water demand, which can be understood and modelled using higher resolution data. While data resolution depends on the installed meter, the logging time can be shortened without installation of smart meters but simply increasing the traditional reading frequency by the users. However, so far only ad-hoc studies systematically collected and analyzed data at daily resolution (e.g., Olmstead et al., 2007; Wong et al., 2010) and few water companies (e.g., Water Corporation in Western Australia and Thames Water in London) started increasing their reading frequency by direct involvement of their customers, who are invited to self-read their consumption and communicate it online to the water company (e.g., Anda et al., 2013).

2.2.2 High resolution data

The advent of high resolution sensors, with their ability of sampling water consumption on sub-daily basis, opened up a new potential to better characterize domestic water consumption. Two distinctive metering approaches can be distinguished: *intrusive metering*, which ensures direct estimates of the residential water end-uses by installing high resolution sensors on-device, namely one sensor for each water consuming appliance (e.g., washing machine, toilet flush, shower-head); *non-intrusive metering*, which registers the total water flow at the household level over one single detection point for the whole house.

Intrusive metering (see Rowlands et al., 2014, and references therein) is generally considered inapplicable in real-world, large-scale analysis as the number of sensors to be installed makes this approach resource intensive, costly, and hardly accepted by household occupants (Cordell et al., 2003; Kim et al., 2008). On the contrary, non-intrusive metering represents a more acceptable, though less accurate, alternative (Mayer and DeOreo, 1999). However, this approach requires disaggregation algorithms to breakdown the total consumption data at the household level into the different end-use categories (see Section 2.3).

Several types of sensors have been developed (Table 2.2) by exploiting different technologies and physical properties of the water flow (for a review see Arregui et al., 2006, and references therein):

- Accelerometers (e.g., Evans et al., 2004), which analyze vibrations in a pipe induced by the turbulence of the water flow. A sampling frequency of 100 Hz of the pipe vibrations allows reconstructing the average flow within the pipe with a resolution of 0.015 liters (Kim et al., 2008).
- Ultrasonic sensors (Mori et al., 2004), which estimate the flow velocity, and then determine the flow rate knowing the pipe section, by measuring the difference in time between ultrasonic beams generated by piezoelectric devices and transmitted within the water flow. The transducers are

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

generally operated in the range 0.5-2 MHz and allow attaining an average resolution around 0.0018 liters (e.g., Sanderson and Yeung, 2002).

- Pressure sensors (Froehlich et al., 2009, 2011), which consist in steel devices, equipped with an analog-digital converter and a micro-controller, continuously sampling pressure with a theoretical maximum resolution of 2 kHz. Flow rate is related to the pressure change generated by the opening/close of the water devices valves via Poiseuille's Law.
- Flow meters (Mayer and DeOreo, 1999), which exploit the water flow to spin either pistons (mechanic flow meters) or magnets (magnetic meters) and correlate the number of revolutions or pulse to the water volume passing through the pipe. Sensing resolution spans between 34.2 and 72 pulses per liter (i.e., 1 pulse every 0.029 and 0.014 liters, respectively) associated to a logging frequency in the range of 1 to 10 seconds (Kowalski and Marshall, 2005; Heinrich, 2007; Willis et al., 2013).

So far, only flow meters and pressure sensors have been employed in *smart meters* applications because ultrasonic sensors are too costly and the use of accelerometers requires an intrusive calibration phase with the placement of multiple meters distributed on the pipe network for each single device of interest (Kim et al., 2008). It is worth noting that the "smartness" of these sensors is related both to their high sampling resolution and to their integration in efficient systems combining data collection, transfer, storage, and analysis. Although sensors can be equipped with data loggers requiring human intervention to retrieve the data directly from the sensors (Mayer et al., 2004), bluetooth and wireless connections have been recently exploited for improving data management. For example, Froehlich et al. (2009) installed a network of pressure sensors communicating via bluetooth with a laptop deployed at each household, which runs a custom data logger to receive, compress, and archive data. These latter are then uploaded to a web server at 30-minute intervals.

2.2.3 Research challenges and future directions

While smart meters are becoming easily available, we identified a list of open research and technical challenges that need to be addressed to promote the coherent use of this wide range of technologies:

1. The first open research question relates to the management of the metered high resolution flow data. In particular, the development of robust, automated processes to transfer the generated big data requires further elaborations, both in terms of hardware and software performance due to existing issues with respect to wireless network reliability, black spots, power source and battery life (Stewart et al., 2010; Little and Flynn, 2012).

Table 2.2: *Studies contributing in the data gathering step. Studies gathering data with a sub-daily resolution are considered as high-resolution, low-resolution otherwise.*

Reference	Location	Resolution	Sensor Type	Resolution [liters]
Olmstead et al. (2007)	USA	low	-	-
Wong et al. (2010)	Hong Kong	low	-	-
Anda et al. (2013)	Australia	low	-	-
Boyle et al. (2013)	N/A	high	-	-
Cordell et al. (2003)	Australia	high	-	-
Kim et al. (2008)	N/A	high	accelerometer	0.0150
Mayer and DeOreo (1999)	USA	high	flow meter	0.014-0.029
Evans et al. (2004)	N/A	high	accelerometer	0.0150
Mori et al. (2004)	N/A	high	ultrasonic	0.0018
Sanderson and Yeung (2002)	N/A	high	ultrasonic	0.0018
Froehlich et al. (2009)	N/A	high	pressure	0.0600
Froehlich et al. (2011)	N/A	high	pressure	0.0600
Kowalski and Marshallsay (2005)	UK	high	flow meter	0.014-0.029
Heinrich (2007)	New Zealand	high	flow meter	0.014-0.029
Willis et al. (2013)	Australia	high	flow meter	0.014-0.029
Mayer et al. (2004)	USA	high	flow meter	0.014-0.029
Mayer et al. (2000)	USA	high	flow meter	0.014-0.029
Mayer et al. (2003)	USA	high	flow meter	0.014-0.029

All these aspects appear key also because the possibility of integrating water and energy meters and using the same data loggers and transmission systems is expected to enhance the diffusion of high resolution water sensors (Benzi et al., 2011; Froes Lima and Portillo Navas, 2012).

2. The second open challenge concerns the design of centralized or distributed information systems to store the data collected by the smart meters (Oracle, 2009). A centralized system would allow checking the accuracy of the collected data, which can then be made easily available for data processing and analysis. On the contrary, a distributed solution would reduce transmission costs and facilitate providing immediate feedbacks to customers, who can use this information to make decisions about their water use.
3. A third open question is how householder privacy is impacted by collection and communication of detailed water-use information. Although such issues are currently underestimated as in many communities (e.g., in Australia) severe water shortages have led to a permissive attitude to conserve water (Giurco et al., 2010), it is likely that the collection of information on both water use and behavior change over time implies increased privacy risks (McIntyre, 2008; Chen et al., 2014).
4. Finally, a challenge is posed by the actual deployment of large-scale high-resolution metering network in the real world. While literature presents

a number of trials (e.g., Mayer et al. (2004); Heinrich (2007); Froehlich et al. (2009)) that exploit smart sensors with extremely fine resolutions (sub-minute), cost, privacy, and regulations may limit their scalability to large-scale continuous operative smart meter installations. For example, data protection and data security issues are being seriously considered by the European Union, which is imposing some strict guidelines to utilities willing to deploy smart meter solutions for their customers and many water utilities collect data at lower resolution than the minute (e.g., Thames Water in London reads data at 15-minute resolution, EMIVASA in Valencia and SES in Switzerland at 1-hour resolution). This implies that the theoretical capabilities of smart metering technologies may not be fully exploited, potentially limiting the accuracy in characterizing the residential water consumption as studies relying on medium/low resolution data. Large-scale smart-meters application would therefore benefit from a better understanding of the consequences of different time resolutions on the models accuracy and on the effectiveness of WDMSSs.

2.3 Water end-uses characterization

Non-intrusive metering requires disaggregation algorithms to breakdown the total consumption data registered at the household level into the different end-use categories (second block of Figure 2.2). In the water research literature, several studies have been conducted in the last two decades using a variety of single or mixed disaggregation methods, such as household auditing, diaries, high resolution flow meters and pressure sensors (see Table 2.3). According to the methodology adopted, we can identify two main approaches for disaggregating smart metered water data at very high temporal resolution: *decision tree algorithms*, namely Trace Wizard[®] (DeOreo et al., 1996) and Identiflow[®] (Kowalski and Marshallsay, 2003), and *machine learning algorithms*, namely HydroSense (Froehlich et al., 2011) and the approach adopted in the SEQREUS project (Beal et al., 2011a). Recently, the disaggregation of medium resolution water data (i.e., hourly data) has been explored by means of water use signature patterns method (Cardell-Oliver, 2013a,b), namely a combination of feature selection, unsupervised learning, and cluster evaluation.

2.3.1 Trace Wizard

Trace Wizard (DeOreo et al., 1996) is a commercial software (recently replaced by an on-demand service developed and managed by Aquacraft Inc) which applies a decision tree algorithm to interpret magnetic metered flow data based on some basic flow boundary conditions (e.g., minimum/maximum volume, peak

flow rate, duration range, etc.). The disaggregation process is structured in the following steps:

1. Conduct a detailed water device stock inventory audit for each household to determine the efficiency rating of each household appliance/fixture;
2. Household occupants should complete a diary of water use events over a one-week period to gain information on their water use habits;
3. Analysts use water audits, diaries, and sample flow trace data for each household to create specific templates that serve to match water end-use patterns depending on some basic flow boundary conditions.
4. Based on the developed templates, stock survey audit, diary information and analysts' experience, the individual water end-uses are disaggregated.

It is worth noting that the human resource effort required by Trace Wizard makes the overall process extremely time and resource intensive, with the quality of the results that is strongly dependent on the experience of the analyst in understanding flow signatures. It has been estimated that the classification of two weeks of data approximatively requires two hours of works by the analyst and attains an average classification accuracy of 70% (Nguyen et al., 2013a). In addition, the prediction accuracy of Trace Wizard is significantly reduced when more than two events occur concurrently (Mayer and DeOreo, 1999). However, Trace Wizard still has an edge on disaggregation techniques and has been used in several research works and projects (DeOreo and Mayer, 1994; Mayer and DeOreo, 1995; DeOreo et al., 1996; Mayer and DeOreo, 1999; DeOreo and Mayer, 2000; Loh et al., 2003; Mayer et al., 2004; Roberts, 2005; Heinrich, 2007; Mead and Aravinthan, 2009; Willis et al., 2009a,b; Aquacraft Inc., 2011; DeOreo et al., 2011).

2.3.2 Identiflow

Similar to Trace Wizard, Identiflow (Kowalski and Marshallsay, 2003) relies on a decision tree algorithm to perform a semi-automatic disaggregation of the total water consumption at the household level. Identiflow uses fixed physical features of various water-use devices (e.g., volume, flow rate, duration, etc.) to classify the different end-use events. Although Identiflow has shown better performance than Trace Wizard (i.e., 74.8% accuracy in terms of the correctly classified volume over 3870 events (Nguyen et al., 2013a)), its classification accuracy strongly depends on the physical features used to describe each fixture/appliance. Two different water events are likely classified into the same category if they exhibit similar physical characteristics. Moreover, it fails to classify events when old devices are replaced by modern ones, since the physical characteristics of these latter might be completely different compared to the old ones.

2.3.3 HydroSense

HydroSense (Froehlich et al., 2011) is a probabilistic-based classification approach which relies on data collected through pressure sensors. Water end-use events are classified with respect to the unique pressure waves that propagate to the sensors when valves are opened or closed. Specifically, when a valve is opened or closed, a pressure change occurs and a pressure wave is generated in the plumbing system. Based on the pressure wave (which depends on the valve type and its location), water end-use events are classified by using advanced pattern matching algorithms and Bayesian probabilistic models. HydroSense has been demonstrated to attain very high levels of classification accuracy, namely 90% and 94% with one or two pressure sensors, respectively (Froehlich et al., 2011). However, the calibration of the algorithm requires an intrusive monitoring period with the installation of a much larger number of pressure sensors connected to each water device (i.e., Froehlich et al. (2011) used 33 sensors in a single household). This requirement significantly constrains the portability of this approach to a wide urban context as it would entail large costs and privacy issues.

2.3.4 SEQREUS project

The end-use disaggregation approach developed in the SEQREUS project (Beal et al., 2011a) proposes a combination of Hidden Markov Models (HMMs), Dynamic Time Warping (DTW), and time-of-day probability to automatically categorize the collected data at the household level into particular water end-use categories. To minimize the intrusiveness of the approach, the ground truth for the calibration (i.e., a set of disaggregated end-use events) is obtained using Trace Wizard. Then, the SEQREUS approach works as follows:

1. The disaggregated data are used for training multiple HMMs, one for each end-use category (excluding the inconclusive event);
2. The physical characteristics of each end-use category are used to refine the estimate given by the HMMs (e.g., any shower event with a volume less than 7 liters or any bathtub event with duration less than 4 minutes is placed in the inconclusive event for future analysis);
3. A DTW algorithm determines if any event in the inconclusive dataset is similar to an event in categories having clearly defined consumption patterns, namely the washing machine and dishwasher cycles;
4. Time of day probability is used to assign inconclusive events to an end-use category.

Testing on three independent households located in Melbourne (Australia) demonstrated a high prediction accuracy, namely between 80% and 90% for the major end-use categories (Nguyen et al., 2014). However, the method still requires human input to achieve such levels of recognition accuracy (e.g., for the classification of inconclusive events supported by DTW and for manually classifying combine events) (Nguyen et al., 2013a,b).

Table 2.3: *Studies contributing in the water end-uses characterization step.*

Reference	Location	Disaggregation algorithm	Number of households
Froehlich et al. (2011)	N/A	HydoSense	5
Heinrich (2007)	New Zeland	Trace Wizard	12
Mayer et al. (2004)	USA	Trace Wizard	33
DeOreo et al. (1996)	USA	Trace Wizard	N/A
Kowalski and Marshallsay (2003)	UK	Identiflow	250
Kowalski and Marshallsay (2005)	UK	Identiflow	N/A
Beal et al. (2011a)	Australia	SEQREUS	1500
DeOreo and Mayer (1994)	USA	Trace Wizard	16
Mayer and DeOreo (1995)	USA	Trace Wizard	16
DeOreo and Mayer (2000)	USA	Trace Wizard	10
Loh et al. (2003)	Australia	Trace Wizard	720
Roberts (2005)	Australia	Trace Wizard	100
Mead and Aravinthan (2009)	Australia	Trace Wizard	10
Willis et al. (2009a)	Australia	Trace Wizard	200
Willis et al. (2009b)	Australia	Trace Wizard	151
Aquacraft Inc. (2011)	USA	Trace Wizard	209
Nguyen et al. (2014)	Australia	SEQREUS	3
Nguyen et al. (2013a)	Australia	SEQREUS	252
Nguyen et al. (2013b)	Australia	SEQREUS	3 (out of 252)
Mayer et al. (2000)	USA	Trace Wizard	37 (out of 1188)
Mayer et al. (2003)	USA	Trace Wizard	33
DeOreo (2011)	USA	Trace Wizard	1000
Cardell-Oliver (2013a)	Australia	Water Use Signature Patterns	11000
Cardell-Oliver (2013b)	Australia	Water Use Signature Patterns	187

2.3.5 Research challenges and future directions

Given the small number of algorithms for disaggregating water flow data, there is still a large room for developing new methods addressing the major limitations of the existing approaches:

1. First, most of the approaches used in the water sector requires time consuming expert manual processing and intensive human interactions via surveys, audits and water event diaries, while the development of automatic procedures is fundamental to further extend the application of these methods beyond experimental trials and research projects (Stewart et al., 2010). Moreover, the existing methods have limited accuracy in identifying overlapping events.

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

The disaggregation problem has been addressed in other research fields as a general problem of *blind identification*, or output-only system identification (Reynders, 2012). The real state of the system (i.e., the set of the working states and water consumption of each single fixture in the household) is unknown and only observations of the system output (i.e., the total water consumption) are available. Starting from the 1990s, several techniques have been proposed to address blind identification problems in different research field, such as signal processing, data communication, speech recognition, image restoration, seismic signal processing (see Abed-Meraim et al., 1997, and references therein).

With the development of smart electricity grids (Kramers et al., 2014; Niesse et al., 2014), this problem has been largely studied in the energy sector to develop automatic disaggregation methods, also known as NILM algorithms, which aim at decomposing the aggregate household energy consumption data collected from a single measurement point into device-level consumption data (for a review, see Zeifman et al., 2011; Zoha et al., 2012; Carrie Armel et al., 2013, and references therein). These methods show promising results and seem effective also up to 6-10 appliances (Figueiredo et al., 2014; Makonin et al., 2013). Yet, the portability of such techniques in the water field has not been assessed. Some additional challenges in characterizing water end-use events might be introduced by the larger human dependency than the one of electric appliances, which are more automatic. These concerns primarily involve manually controlled fixtures (e.g., bathtubs, showers, faucets), which might be used not at the maximum capacity (Froehlich et al., 2009).

2. The second main open question relates to the acquisition of the ground truth for initial calibration. All the algorithms used for disaggregating water data, but also the majority of the ones used for energy data, need an intrusive period to collect a dataset of disaggregated end-use events, which incurs extra cost and human effort, ultimately challenging their large-scale application. Researchers are actively looking to devise completely unsupervised or semi-supervised methods that avoid the effort of acquiring the calibration ground truth data (e.g., Goncalves et al., 2011; Parson et al., 2014).
3. Finally, most of the approaches developed in the energy sector are currently focused on correctly characterizing the on/off status of the devices and, possibly, the fraction of total energy assigned correctly, while their performance in reproducing the timings and frequencies of each device are lower (Batra et al., 2014). Yet, timings and frequencies represent key information to understand consumers behaviors and design personalized

demand management strategies (e.g., deferring the use of some appliances to peak-off hours). Accordingly, knowledge about use frequencies, timing and peak-hours in the water sector would constitute crucial information for identifying both typical consumption behaviours and patterns, as well as consumption anomalies (e.g., leakages (Loureiro et al., 2014; Ponce et al., 2014; Pérez et al., 2014; Perez et al., 2014)). This knowledge would aid the activities of water utilities at different levels: demand management, network maintenance, and strategic planning.

2.4 User modeling

The user modeling phase (third block in Figure 2.2) aims at representing the water demand at the household level, thus preserving the heterogeneity of the individual users in the modelled community, possibly as determined by natural and socio-psychographic factors as well as by the users' response to different WDMs. In the literature, two distinctive approaches exist (see Table 2.4): *descriptive models*, which limit their extent to the analysis of water consumption patterns, and *predictive models*, which provide estimate of the water consumption at the individual (household) level as determined by natural and socio-psychographic factors, and in response to different WDMs.

2.4.1 Descriptive models

The first class of models, namely descriptive models, aims at analyzing the observed water consumption behaviors of water users. Depending on the resolution of the data available, the analysis can focus on identifying aggregated consumption patterns or on defining users' profiles on the basis of the disaggregated end-uses (e.g., Loh et al., 2003; SDU, 2011; SJESD, 2011; Gato-Trinidad et al., 2011; Willis et al., 2011; Beal et al., 2011b, 2013; Cardell-Oliver and Peach, 2013; Cole and Stewart, 2013; Beal and Stewart, 2014; Beal et al., 2014; Gurung et al., 2014, 2015).

The construction of descriptive models allows studying historical trends (Agudelo-Vera et al., 2014; Kofinas et al., 2014) to build a user consumption profile that constitutes the baseline for identifying the most promising areas where conservation efforts may be polarized (e.g., restriction on irrigation practices in case gardening represents the dominant end-use). However, the majority of these models cannot be used to predict the water savings potential of alternative WDMs, unless combined with control group experiments to observe user responses (Cahill et al., 2013).

2.4.2 Predictive models

The second class of models, namely predictive models, aims at estimating the water demand at the individual (household) level. Some works developed predictive models that mostly provide short-term forecast of the water demand on the basis of time series analyses (e.g., Homwongs et al., 1994; Molino et al., 1996; Altunkaynak et al., 2005; Alvisi et al., 2007; Nasserri et al., 2011). Yet, these approaches are ineffective in supporting the design and implementation of WDMSs as the predicted water consumption of a user is not related to his socio-psychographic factors or his response to different WDMSs. An alternative approach can be structured in the following two sub-steps: (i) *multivariate analysis*, which consists in the identification and selection of the most relevant inputs to explain the preselected output, and (ii) *behavioral modeling*, which means model structure identification, parameter calibration and validation.

The multivariate analysis phase (i.e., variable selection as called in data-driven modeling (George, 2000)) is a fundamental step to build predictive models of urban water demand variability in space and time. In most of the works, the identification of the most relevant drivers relies on the results of data mining techniques (e.g., correlation analysis) between a pre-defined set of variables (candidate drivers) and the water consumption data. This approach is also referred to as *inductive* modelling (Cahill et al., 2013). An alternative to this data-driven approach is the *deductive* construction of models according to empirical or theoretical causality (Cahill et al., 2013). Depending on the specific domains from which the candidate drivers are extracted, which is often delimited by data availability (Arbués et al., 2003), we can distinguish the following three main approaches:

- *economic-driven studies*, which focus on studying the correlation between water consumption and purely economic drivers, such as water tariff structures or water price elasticity (e.g., Schneider and Whitlatch, 1991; Espey et al., 1997; Brookshire et al., 2002; Dalhuisen et al., 2003; Olmstead et al., 2007; Olmstead and Stavins, 2009; Rosenberg, 2010; Qi and Chang, 2011);
- *geo-spatial studies*, which assess the correlation between hydro-climatic variables and seasonality with water consumption (e.g., Miaou, 1990; Griffin and Chang, 1991; Zhou et al., 2000, 2002; Fullerton and Elias, 2004; Aly and Wanakule, 2004; Gato et al., 2007; Balling and Gober, 2007; Balling et al., 2008; Lee and Wentz, 2008; Praskievicz and Chang, 2009; Corbella and Pujol, 2009; Chang et al., 2010; Polebitski and Palmer, 2010; Lee and Wentz, 2010; Lee et al., 2011);
- *psycographic-driven studies*, which infer the influence of users' personal attributes on their water consumption, including income, family composition, lifestyle, and households physical characteristics (e.g., number of

rooms, type, presence of garden) (e.g., Syme et al., 2004; Wentz and Gober, 2007; Fox et al., 2009; Jorgensen et al., 2009; Russell and Fielding, 2010; Grafton et al., 2011; Willis et al., 2013; Suero et al., 2012; Matos et al., 2014; Talebpour et al., 2014; Romano et al., 2014).

Note that this classification is not stringent, in the sense that hybrid approaches dealing with more than one of the mentioned domains have already been developed (e.g., Makki et al., 2015). Similarly to the descriptive models discussed in the previous section, the development of predictive models could significantly benefit from smart metering technologies and high-resolution water consumption data. Indeed, the availability of high-resolution and end-use characterization of the water consumption allows predicting the effects of customized WDMSs focused on specific end-uses (e.g., Makki et al., 2013). In most of the literature, the user modeling is limited to the multivariate analysis, which however provides only qualitative information to water managers, water utilities, and decision makers. Only few works completed the second phase (i.e., behavioral modeling) and provide a quantitative prediction of the water demand at the household level, thus representing better decision-aiding tools as they can use these models to develop what-if analysis as well as scenario simulation and analysis.

The construction of behavioral models aims at the identification, calibration, and validation of mathematical models, which describe the water demand (i.e., output variable) as a function of the drivers identified in the multivariate analysis. In the behavioral modeling literature, we can identify a first class of models, named *single-user models*, which describe the consumption behavior of individual users considered as isolated entities. These works (e.g., Lyman, 1992; Gato, 2006; Kenney et al., 2008; Maggioni, 2015) generally rely on dynamic models based on sampling of statistical distributions describing average users and end-uses (e.g., number of people per household and their ages, the frequency of use, flow duration and event occurrence likelihood). Water demand patterns can be then estimated via model simulation and comparison of the results with the observed data. Yet, this approach often reduces the heterogeneity of the water users, which can be preserved by running Monte Carlo simulations that sample also the extreme values of the associated statistical distributions (Rosenberg et al., 2007; Blokker et al., 2010; Cahill et al., 2013). Recently, different approaches (Bennett et al., 2013; Makki et al., 2013, 2015) combining non-parametric statistical tests and advanced regression models to identify key water consumption drivers and forecast urban water consumption have been demonstrated to successfully identify the main drivers of water consumption and to attain good forecast accuracy levels.

A second class of behavioral models, named *multi-user models*, instead focus on studying the social interactions and influence/mimicking mechanisms

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

among the users. The majority of these works relies on multiagent systems (Shoham and Leyton-Brown, 2009), where each water user (agent) is defined as a computer system situated in some environment and capable of autonomous actions to meet its design objectives, but also able to exchange information with the neighbor agents and change its behavior accordingly (Wooldridge, 2009). The adoption of agent-based modeling offers several advantages with respect to other approaches (Bonabeau, 2002; Bousquet and Le Page, 2004): (1) it provides a more natural description of a system, especially when it is composed of multiple, distributed, and autonomous agents, (2) it relaxes the hypothesis of homogeneity in a population of actually heterogeneous individuals, (3) it allows an explicit representation of spatial variability, and (4) it captures emergent global behaviors resulting from local interactions. As a consequence, multiagent systems can be employed to study the role of social network structures and mechanisms of mutual interaction and mimicking on the behaviors of water users (e.g., Rixon et al., 2007; Galán et al., 2009), to estimate market penetration of water-saving technologies (e.g., Chu et al., 2009), and to simulate the feedbacks between water consumers and policy makers (e.g., Kanta and Zechman, 2014).

Table 2.4: *Studies contributing in the user modeling step. Legend for multivariate analysis approaches: E = economic-driven; GS = geo-spatial; P = psychographic driven; AR = autoregressive. Legend for behavioural models approach: single = single user model; multi = multi-user model.*

Reference	Location	Modeling approach	Multivariate analysis	Behavioural model	Spatial scale
Loh et al. (2003)	Australia	descriptive	-	-	household
Gato-Trinidad et al. (2011)	Australia	descriptive	-	-	household
SDU (2011)	USA	descriptive	-	-	household
SJESD (2011)	USA	descriptive	-	-	household
Cardell-Oliver and Peach (2013)	Australia	descriptive	-	-	household
Beal et al. (2013)	Australia	descriptive	-	-	household
Beal and Stewart (2014)	Australia	descriptive	-	-	household
Gurung et al. (2015)	Australia	descriptive	-	-	household
Gurung et al. (2014)	Australia	descriptive	-	-	household
Beal et al. (2014)	Australia	descriptive	-	-	household
Cole and Stewart (2013)	Australia	descriptive	-	-	household
Willis et al. (2011)	Australia	descriptive	-	-	household
Beal et al. (2011b)	Australia	descriptive	-	-	household
Maggioni (2015)	USA	predictive	E+GS+P	single	household
Makki et al. (2015)	Australia	predictive	E+P	single	household
House-Peters and Chang (2011)	N/A	predictive	E+GS+P	single+multi	N/A
Schneider and Whitlatch (1991)	USA	predictive	E	-	district
Lyman (1992)	USA	predictive	E+GS+P	single	household
Espey et al. (1997)	N/A	predictive	E	-	N/A
Dalhuisen et al. (2003)	N/A	predictive	E	-	N/A
Miaou (1990)	USA	predictive	GS	-	urban
Polebitski and Palmer (2010)	USA	predictive	GS	-	census tracts
Lee et al. (2011)	USA	predictive	GS	-	household
Olmstead et al. (2007)	USA	predictive	E	-	household
Willis et al. (2013)	Australia	predictive	P	-	household
Homwongs et al. (1994)	USA	predictive	AR	-	urban
Molino et al. (1996)	Italy	predictive	AR	-	urban
Altunkaynak et al. (2005)	Turkey	predictive	AR	-	urban
Alvisi et al. (2007)	Italy	predictive	AR	-	household
Nasseri et al. (2011)	Iran	predictive	AR	-	urban
Brookshire et al. (2002)	N/A	predictive	E	-	N/A
Olmstead and Stavins (2009)	N/A	predictive	E	-	N/A
Rosenberg (2010)	Jordan	predictive	E	-	household
Qi and Chang (2011)	USA	predictive	E	-	urban

Table 2.4: (Continued) Studies contributing in the user modeling step.

Reference	Location	Modeling approach	Multivariate analysis	Behavioural model	Spatial scale
Griffin and Chang (1991)	USA	predictive	GS	-	district
Zhou et al. (2000)	Australia	predictive	GS	-	urban
Zhou et al. (2002)	Australia	predictive	GS	-	district
Fullerton and Elias (2004)	USA	predictive	GS	-	urban
Aly and Wanakule (2004)	USA	predictive	GS	-	urban
Gato et al. (2007)	Australia	predictive	GS	-	urban
Balling and Gober (2007)	USA	predictive	GS	-	urban
Balling et al. (2008)	USA	predictive	GS	-	census tracts
Lee and Wentz (2008)	USA	predictive	GS	-	census tracts
Praskievicz and Chang (2009)	Korea	predictive	GS	-	urban
Corbella and Pujol (2009)	N/A	predictive	GS	-	N/A
Chang et al. (2010)	USA	predictive	GS	-	household
Lee and Wentz (2010)	USA	predictive	GS	-	urban
Syme et al. (2004)	Australia	predictive	P	-	household
Wentz and Gober (2007)	USA	predictive	P	-	household
Fox et al. (2009)	UK	predictive	P	-	household
Russell and Fielding (2010)	N/A	predictive	P	-	N/A
Grafton et al. (2011)	10 OECD countries	predictive	P	-	household
Suero et al. (2012)	USA	predictive	P	-	household
Matos et al. (2014)	Portugal	predictive	P	-	household
Talebpour et al. (2014)	Australia	predictive	P	-	household
Romano et al. (2014)	Italy	predictive	P	-	water utility
Gato (2006)	Australia	predictive	GS	single	urban
Rosenberg et al. (2007)	Jordan	predictive	GS+P	single	household
Blokker et al. (2010)	Nederland	predictive	P	single	household
Cahill et al. (2013)	USA	predictive	P	single	household
Bennett et al. (2013)	Australia	predictive	GS+E+P	single	household
Rixon et al. (2007)	Australia	predictive	E+P	multi	household
Galán et al. (2009)	Spain	predictive	P	multi	household
Chu et al. (2009)	China	predictive	E+P	multi	household
Kanta and Zechman (2014)	N/A	predictive	GS+P	multi	household
Jorgensen et al. (2009)	N/A	predictive	P	-	household
Kenney et al. (2008)	USA	predictive	E+GS+P	single	household
Makki et al. (2013)	Australia	predictive	E+P	single	household
Giacomoni and Berglund (2015)	USA	predictive	GS	multi	urban
Escriva-Bou et al. (2015b)	USA	predictive	P	single	household
Escriva-Bou et al. (2015a)	USA	predictive	P	single	household

2.4.3 Research challenges and future directions

Given the current status of user modeling studies and the room for improvement given by the use of high resolution, smart metered data, several research challenges and future directions emerge:

1. The first open question in terms of descriptive models concerns matching the analysis of the water consumption patterns with the potential drivers generating the observed users' behaviors. This would allow validating the results of the classification of the users on the basis of their consumption and understanding if this latter is a good proxy representing different characteristics of the users.
2. The use of spatially explicit models to take advantage of the high temporal and spatial resolution of smart metered data is often hindered by the aggregation of individual household data to a larger spatial scale to protect customers' privacy as well as by the difficulties in collecting and sharing data coming across multiple water authorities and administrative institutions (House-Peters and Chang, 2011).

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

3. The third major challenge relates to the validation of the agent-based behavioral models. As in the construction of complex process-based models, accurately describing the single user (agent) behavior and connecting multiple users within an agent-based model does not ensure the validity of the results, although these latter are contrasted with observed data. In addition, given the large number of assumption and parameters, the problem of equifinality (i.e., the potential existence of multiple, alternative parameterization leading to same simulation outcomes) has to be addressed (Ligtenberg et al., 2010).
4. It is worth noting that the type of candidate drivers considered in the user modeling phase impacts the statistical representativeness of the results. The construction of sufficiently large datasets to estimate the relationships between water consumption data and the uncontrolled drivers (i.e., hydro-climatic and psychographic variables) is generally easy, provided that the time period is long enough and the number of involved users is sufficiently high. On the contrary, in most of the cases there is a single historical realization of the controllable drivers, namely the ones subject to human decisions (e.g., the existing pricing scheme). In such cases, the response of the users to different options is generally estimated via economics principles or surveys. Yet, economic principles introduce a priori general rules that might be inaccurate in characterizing the specific users under study, and the surveys provide only a static snapshot of the system conditions. The potential for using experimental trials (e.g., Gilg and Barr, 2006; Borisova and Useche, 2013; Fielding et al., 2013) and gamification platforms (e.g., Mühlhäuser et al., 2008) to validate behavioral models results by retrieving information to the real users in large-scale applications has not been tested yet.
5. Finally, a major opportunity is represented by the development of integrated models that cross-analyze water and water-related energy consumption data to improve residential water demand models (Abdallah and Rosenberg, 2014; Escriva-Bou et al., 2015b,a).

2.5 Personalized water demand management strategies

Literature reports of a variety of management policies acting on the demand side of residential water consumption, designed with the purpose of improving water conservation and safeguarding water security in urban contexts.

According to Inman and Jeffrey (2006), they can be classified in the following five categories (Table 2.5): *technological*, *financial*, *legislative*, *maintenance*, and *educational*. These strategies differ in the time scales they act on: price and pre-

2.5. Personalized water demand management strategies

scriptive (i.e., command-and-control) approaches have been shown to achieve significant reductions of water demand in the short-period, but also have some drawbacks (such as equity issues and limits in consumers' price elasticity) that may limit the effectiveness of such strategies in the long term, if not integrated with other water conservation interventions (Fielding et al., 2013; Renwick and Green, 2000). In contrast, users' awareness and educational approaches allow for smaller reductions in the short period, but appear to be crucial to pursue reductions on the long run, as they require a change in users' behaviors (Geller, 2002).

Technological strategies involve the installation of water efficient household appliances (e.g., Mead and Aravinthan, 2009; Suero et al., 2012; Carragher et al., 2012; Froes Lima and Portillo Navas, 2012; Gurung et al., 2015). This option offers great potential for reducing indoor and outdoor water consumption (Mayer et al., 2000, 2003, 2004; DeOreo, 2011). Yet, the benefits associated to these advanced systems are inconstant (Maggioni, 2015). For example, an incorrect use of automatic sprinkler may consume more water than manually operated irrigation systems (Syme et al., 2004), thus requiring educational programs to ensure an appropriate use.

Financial strategies, (also called market-based or price approaches (Olmstead and Stavins, 2009)), consist in water tariffs control associated to analysis of water demand elasticity (e.g., Dandy et al., 1997; Dalhuisen et al., 2003; Arbués et al., 2003; Kenney et al., 2008; Cole et al., 2012; Molinos-Senante, 2014; Maggioni, 2015). Even though some authors claim that price-based strategies are more cost effective than other conservation programs (Olmstead and Stavins, 2009), the effectiveness of this strategies seems uncertain as water demand has been shown to be relatively price inelastic (Worthington and Hoffman, 2008) and to rebound to the same or even higher levels after an initial decrease (Kanakoudis, 2002). Yet, a careful assessment of the effectiveness of these strategies would benefit from longer dataset gathered in multiple jurisdictions and contexts (Worthington and Hoffman, 2008). In addition, there are also concerns about the equity of raising prices (Duke et al., 2002).

Legislative strategies correspond to mandatory regulations and restrictions on water use, particularly in case of drought (e.g., Kenney et al., 2004; Hensher et al., 2006; Brennan et al., 2007; Kenney et al., 2008; Grafton and Ward, 2008). Restrictions applied to specific water uses, such as car washing or irrigation, have been demonstrated to reduce water consumption up to 30% (Renwick and Archibald, 1998; Kanakoudis, 2002). However, they require policy intervention to be implemented (Maggioni, 2015) and may be resisted by the community (Steg and Vlek, 2009).

Maintenance strategies consist in operations aiming at reducing or eliminating leakages in the water supply networks (e.g., Britton et al., 2008, 2013), which

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

generally account for a significant fraction of the water consumption (e.g., EEA (2001) estimated losses due to leakage equal to 30% in Italy and 50% in Bulgaria). The identification and repair of leakages, which are often associated to a small number of households (Roberts, 2005; Mayer and DeOreo, 1999; Mayer et al., 2004), allows substantial increase in the efficiency of the water supply systems at lower costs with respect to augmenting the water supplied without repairing the network (Garcia and Thomas, 2001; Brooks, 2006).

Educational strategies aim at engaging the water users by means of public awareness and education campaigns (e.g., Geller, 2002; Steg and Vlek, 2009; Froes Lima and Portillo Navas, 2012; Anda et al., 2013; Fielding et al., 2013; Stewart et al., 2013). The effectiveness of these approaches is case-dependent: for example, it is estimated that information campaigns successfully led to a reduction of water demand equal to 8% in the period 1989-1996 in California (Renwick and Green, 2000), while no impact was observed in UK, where, although a large campaign involving direct mailing as well as newspaper and radio advertisements, only 5% of the 8000 residences involved noticed the campaign (Howarth and Butler, 2004). Recent studies however suggest that a relevant water saving potential can be obtained by providing feedbacks to the users about their water consumption or suggestions on customized water savings practices (e.g., Kenney et al., 2008; Willis et al., 2010; Froehlich et al., 2012; Sonderlund et al., 2014).

Regardless the type of demand-side management strategy implemented, the availability of high-resolution data appears crucial both for the design and for an accurate evaluation of the effects of such interventions. Studies like Mayer et al. (2000) and Mayer et al. (2003), for instance, demonstrate that smart metered data and end-use characterization are crucial tools for evaluating the effects of retrofitting interventions both in terms of consumption reduction for particular end-uses and changes in consumption patterns (i.e., use frequencies and volumes). The same stands for price-based approaches, as smart metered data can be exploited to differentiate the price elasticity in relation to different uses (e.g., outdoor and indoor water consumption), allowing for the design of new price schemes, such as *Time of Use Tariffs* (Cole et al., 2012). In turn, if we consider educational campaigns, there is evidence of the potential of high-resolution metering in supporting the design of effective feedbacks and assess behavioural changes (Froehlich et al., 2012; Stewart et al., 2013; Sonderlund et al., 2014).

2.5.1 Research challenges and future directions

Given the recent improvements in characterizing water users' behaviors, a list of open research challenges exists to improve the designed of personalized WDMSs:

2.5. Personalized water demand management strategies

Table 2.5: *Studies contributing in the personalized WDMSs step. Different WDMSs are considered: E = educational; F = financial; L = legislative; M = maintenance; T = technological.*

Reference	Location	Type of WDMS	Personalized
Maggioni (2015)	USA	L+T+F	x
Inman and Jeffrey (2006)	N/A	T+F+L+M+E	
Britton et al. (2008)	Australia	M	x
Dalhuisen et al. (2003)	N/A	E	
Mayer and DeOreo (1999)	USA	M	x
Mayer et al. (2004)	USA	T+M	x
Roberts (2005)	Australia	M	x
Suero et al. (2012)	USA	T	x
Mayer et al. (2000)	USA	T	x
Mayer et al. (2003)	USA	T	x
DeOreo (2011)	USA	T	x
Dandy et al. (1997)	Australia	F	
Arbués et al. (2003)	N/A	F	
Molinos-Senante (2014)	Spain	F	
Worthington and Hoffman (2008)	N/A	F	
Kanakoudis (2002)	Greece	F	
Duke et al. (2002)	USA	F	
Hensher et al. (2006)	Australia	L	x
Brennan et al. (2007)	Australia	L	
Grafton and Ward (2008)	Australia	L	
Renwick and Archibald (1998)	USA	L	x
Steg and Vlek (2009)	N/A	L-E	x
Britton et al. (2013)	Australia	M	x
Garcia and Thomas (2001)	France	M	
Brooks (2006)	N/A	M	
Fielding et al. (2013)	Australia	E	x
Renwick and Green (2000)	USA	E	
Howarth and Butler (2004)	UK	E	x
Geller (2002)	N/A	E	x
Willis et al. (2010)	Australia	E	x
Froehlich et al. (2012)	USA	E	x
Sonderlund et al. (2014)	N/A	E	x
Kenney et al. (2004)	USA	L	
Kenney et al. (2008)	USA	L+F+E	x
Mead and Aravinthan (2009)	Australia	T	x
Froes Lima and Portillo Navas (2012)	Brazil	T+E	x
Carragher et al. (2012)	Australia	T	x
Cole et al. (2012)	Australia	F	x
Stewart et al. (2013)	Australia	E	x
Gurung et al. (2015)	Australia	T	x
Giacomoni and Berglund (2015)	USA	L+T	
Escriva-Bou et al. (2015a)	USA	T+E	
Escriva-Bou et al. (2015b)	USA	T+E	

1. The first challenge is the identification of more effective strategies for influencing the users behaviors. Technological strategies mostly impact on a limited number of end-uses (e.g., clothes or dish washers), whereas are less effective in inducing water savings in more human-controlled end-uses,

such as showering or tap water. Moreover, investment inefficiencies can limit the effectiveness of these strategies causing the *Efficiency Gap* that is well-known in the energy field (Allcott and Greenstone, 2012). Educational intervention and programs can be more effective in controlling these latter, for example by providing feedbacks to the users as already applied in the energy sector (e.g., Abrahamse et al., 2007; Costanza et al., 2012). Yet, there are still open questions on the use of feedbacks to reduce water (or energy) consumption, particularly with respect to the most effective feedback format, whether the effect persists over time, as well as assessments of costs and benefits of feedback (Strengers, 2011; Desley et al., 2013).

2. The second main open question relates to the long-term effect of WDMSs, especially for educational programs and awareness campaigns (e.g., Peschiera et al., 2010; Pereira et al., 2013). Although they showed promising results during the program and some months afterwards, their effect eventually dissipated and water consumption returned to pre-intervention levels after approximately 12 months (Fielding et al., 2013).
3. Finally, further effort should be devoted to examine the role of social norms and social influence in promoting water conservation (Rixon et al., 2007; Van Der Linden, 2013; Schultz et al., 2014). In particular, the potential for using gamification platforms and social applications to allow users monitoring their consumption coupled with normative information about similar households in their neighborhood should be assessed (Bogost, 2007; Rizzoli et al., 2014; Harou et al., 2014; Clifford et al., 2014; Curry et al., 2014; Savić et al., 2014; Vieira et al., 2014; Kossieris et al., 2014; Magiera and Froelich, 2014; Laspidou, 2014). Water utilities can indeed take advantage of people's tendency to mimic the behavior of their neighbors in order to target their efforts to "early adopters" and encourage technology diffusion (Janmaat, 2013).

2.6 Discussion

In this chapter, we reviewed 134 papers (Table 2.1) that contributed new methodologies and tools in one or more of the blocks underlying the general 4-step procedure represented in Figure 2.2.

A "roadmap" of the main research challenges that need to be addressed in order to move the application of smart meters forward over the next decade is shown in Table 2.6 and summarized below:

1. Data gathering: (i) how to efficiently and reliably manage the big data generated by the acquisition of high resolution smart metered flow data;

- (ii) understanding the best information system architecture (i.e., centralized or distributed) to store the data collected by the smart meters; (iii) how householder privacy is impacted by collection and communication of detailed water-use information;
2. Water End-uses characterization: (i) development of automatic procedures for disaggregating water consumption data at the household level to reduce the manual processing and intensive human interactions required by current methods; (ii) development of unsupervised methods that avoid the effort of acquiring the ground truth for training the algorithms; (iii) enhancing the accuracy of the methods in reproducing the timings and frequencies of each device usage.
 3. User modeling: (i) matching the analysis of the observed water consumption profiles identified in the descriptive models with the potential drivers generating the observed users' behaviors; (ii) better exploit the high spatial resolution of smart metered data to identify water use patterns across geographic areas; (iii) validation of the agent-based behavioral models' simulation against observed data; (iv) testing of experimental trials and gamification platforms to support the validation of the behavioral models as well as to retrieve information from the water users; (v) developing integrated models for water and water-related energy.
 4. Personalized water demand management strategies: (i) identification of more effective strategies for influencing the users behaviors, particularly by means of customized feedbacks to the water users providing information about their water consumption or suggestions on water savings practices; (ii) how to ensure a long-term effect of the implemented water demand management strategies, especially for educational programs and awareness campaigns; (iii) a better understanding of the role of social norms and social influence in promoting water conservation;

Despite the large number of papers published over the last years, the analysis of the studies discussed in this review highlights a clear need to shift research efforts from the development of specialized methodologies within each step of the procedure toward a more integrated approach that covers all the four phases. Indeed, the majority of the studies reviewed (i.e., 89% over 134 papers) provides contribution to a single step, whereas only few works go across multiple steps.

Moreover, we can observe that the case study locations are not homogeneously distributed: 79% of the papers reviewed are applied in the United States (36%) or Australia (43%), while the remaining studies were developed in Europe (13%) or Asia (6%) and a single application found in South America and no one in Africa. However, we expect that the challenges posed by climate

2. Using smart meters for advancing residential water demand modeling and management: state-of-the-art review

Table 2.6: *Main research challenges for the use of smart meters in residential water demand modeling and management.*

1) Data gathering	2) Water end-uses characterization	3) User modeling	4) Personalized WDMSs
1.1) Management of big data	2.1) Automatic disaggregation procedures (i.e., no manual processing)	3.1) Matching observed water consumption profiles with potential drivers of users' behaviors	4.1) More effective behavioral influence via customized feedbacks
1.2) Centralized or distributed information system	2.2) Unsupervised disaggregation algorithms (i.e., no ground truth)	3.2) Identification of spatial patterns across geographical areas	4.2) Long-term effect of WDMS
1.3) Impacts on household privacy	2.3) Higher accuracy in reproducing timings and frequencies	3.3) Validation of the agent-based behavioral models	4.3) Social norms and social influence
1.4) Real world scalability of high-resolution networks		3.4) Testing experimental trials and gamification platforms	
		3.5) Developing integrated models for water and water-related energy	

change impacts, growing population demands, and constrained sources of water supply will call for the application of integrated residential water demand modeling and management in several countries across the world. Finally, we foresee that the investments for smart technologies in fields other than urban water management (e.g., Fernández et al., 2014; Niesse et al., 2014; Kramers et al., 2014; Rezgui et al., 2014; Zarli et al., 2014) will create opportunities for collaborations and common actions among different spheres. Residential water demand modelling and management can benefit from these collaborations because smart technologies and networks have already been deployed in other fields, like domestic energy, thus representing a benchmark for learning and integration. Moreover, the existing nexus between energy and water is expected to foster synergies and cross-influences for addressing future demands (WWAP, 2014; Escriva-Bou et al., 2015b). Integrated, interdisciplinary science will thus support policy makers and planners addressing the major sustainability challenges placed by modern urban contexts and their evolution towards smart cities (Hilty et al., 2006; Laniak et al., 2013; Letcher et al., 2013).

Starting from these remarks, in the next chapters of this thesis we present a set of modelling tools that we developed for each phase of the procedure of Figure 2.2. Addressing various aspects of residential water (and energy) demand modelling and management, we aim to provide relevant scientific contributions towards the advancement of solutions to the challenges above listed.

3

Water and energy end-use disaggregation

Abstract¹

Information on residential water and power consumption patterns disaggregated at the single-appliance level is an essential requirement for water and energy utilities and managers to design customized demand management strategies. Several NILM techniques have been proposed in the literature for power load disaggregation. They decompose the aggregated electric load measured at the household level by a single-point smart meter into the individual contribution of each end-use. Despite being defined *non-intrusive*, NILM methods often require an intrusive data sampling process for training purpose. This calibration intrusiveness hampers NILM methods large-scale applications. Other NILM challenges are the limited accuracy in reproducing the end-use consumption patterns and their trajectories in time, which are key to characterize con-

¹The content of this chapter has been partially published in the following papers:

- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2017. A hybrid signature-based iterative disaggregation algorithm for non-intrusive load monitoring. *Applied Energy* 185, 331–344;
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A. E., 2016. Sparse optimization for automated energy end use disaggregation. *IEEE Transactions on Control Systems Technology* 24 (3), 1044–1051;
- Cominola, A., Giuliani, M., Castelletti, A., Abdallah, A., Rosenberg, D., 11 – 14 July 2016a. Developing a stochastic simulation model for the generation of residential water end-use demand time series. In: *Proceedings of the 8th International Congress on Environmental Modeling and Software (iEMSs 2016)*. Toulouse (France);
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A., 28 June – 3 July 2015. A convex optimization approach for automated water and energy end use disaggregation. In: *Proceedings of the 36th IAHR World Congress*. Vol. 28. The Hague (the Netherlands).

3. Water and energy end-use disaggregation

sumers' behaviours and appliances efficiency, and the poor performance when multiple appliances are simultaneously operated. End-use disaggregation in the water research literature is more recent and fewer algorithms for water end-use disaggregation have been proposed so far. Yet, further developments to compare existing techniques and assess the portability of power load disaggregation algorithm to water data area needed. In this chapter we contribute two algorithms for water and energy end-use disaggregation. The first algorithm is based on the assumption that the unknown appliance power/water consumption profiles are piecewise constant over time and it exploits the information on the time-of-day probability in which a specific appliance might be used. The disaggregation problem is formulated as a least-square error minimization problem, with an additional (convex) penalty term aiming at enforcing the disaggregate signals to be piecewise constant over the time. The second algorithm is called HSID and is based on the combination of FHMMs, which provide an initial approximation of the end-use trajectories, and ISDTW, that processes the end-use trajectories in order to match the typical power/water consumption pattern of each appliance. Tests onto real-world power consumption data benchmarking against state-of-the-art algorithms show that our algorithms achieve high disaggregation accuracy with multiple data resolutions. HSID is also demonstrated to be robust with respect to noisy signals, scalable to dataset including a large set of appliances and can be successfully used in non-intrusive experiments without requiring appliance-level measurements. Finally, we extend the application of HSID to the disaggregation of synthetic water data at multiple data resolutions.

3.1 Introduction

In Chapter 2 (Section 2.3) we motivated how demand management benefits of end-use disaggregation, both in the water and energy fields. As mentioned, pioneering research and applications on end-use disaggregation regard the energy sector: for this reason most of the background theory and literature we mention in this chapter, prior to applications, refer to energy end-use disaggregation, rather than water. Indeed, the effectiveness of customized energy consumption feedbacks and, broadly, demand management strategies, such as economic incentives to upgrade poorly efficient energy consuming devices (Geller et al., 2006), hourly dynamic energy pricing to reduce demand in peak hours (Gaiser and Stroeve, 2014), and awareness campaigns to inform energy consumers about their broken-down consumption and savings (Vassileva and Campillo, 2014), has been demonstrated to benefit from appliance-specific information since the beginning of 1990s (Newborough and Probert, 1990; Fischer, 2008). The knowledge of timings, peak-hours, and frequencies of use of electric de-

vices is key to understand consumers' behaviours, identify consumption anomalies, and, ultimately, design personalized demand management strategies (e.g., deferring the use of some appliances to peak-off hours). Appliance-specific personalized recommendations are potentially worth more than 12% reduction in annual domestic consumption and can bring multiple benefits to energy consumers, utilities, and research and development centres (Carrie Armel et al., 2013). This has been prompting big investments for the deployment of smart metering networks (Neenan and Hemphill, 2008; Chou and Yutami, 2014; Colak et al., 2015), along with the development of *Non-Intrusive Load Monitoring* NILM techniques.

The main advantage of NILM (Hart, 1992) is that it allows decomposing the aggregated electric load measured at the household level by a single smart, high-frequency, meter into the individual contribution by each appliance, the so-called end-uses. Despite alternative options do exist for monitoring residential energy consumption at the appliance level (e.g., smart appliances, distributed sensing networks for direct measurement and smart plugs (Kobus et al., 2015; Morsali et al., 2012)), NILM methods, coupled with single-point sensors, are so far the most promising decomposition approach as they reduce hardware costs (sensor cost and related costs for installation, maintenance, battery and sensors replacement) as well as intrusiveness into users' houses, even though many require an intrusive calibration phase. Also, installing a unique high-resolution sensor per house significantly reduces the amount of data to manage, rather than collecting records from multiple sensors. Another reason promoting the suitability of NILM methods for large-scale energy disaggregation applications and market penetration consists in the overall economic advantages of disaggregation software technologies: a business case by Carrie Armel et al. (2013) shows that the benefits per kWh in terms of potentially avoided energy generation and distribution outweigh the costs of disaggregation technologies by a factor of four. This is further demonstrated by the fact that NILM methods are currently used in domains other than energy consumption, including water and gas, and many companies such as General Electric, Opower and Belkin are working on their development closely with smart meter producers (Zoha et al., 2012; Carrie Armel et al., 2013). Yet, the problem of disaggregating an electric signal into its sub-components places a twofold challenge. On the one hand, disaggregation techniques should be able to maximize the appliance-specific information extracted from the aggregate signal. On the other hand, the algorithms should allow for scalability, while minimizing economic and privacy costs related to disaggregation activities (i.e., sensors installation, data collection, and data analysis).

Several NILM algorithms have been proposed in the literature (see Zeifman et al., 2011; Zoha et al., 2012, and references therein for a review). Yet, a number

3. Water and energy end-use disaggregation

of research and operational challenges are under debate and emerged in recent works. The first, most important issue is related to the rate of intrusiveness of the data sampling process (Goncalves et al., 2011). In fact, after the seminal work by Hart (1992), a first class of supervised algorithms has been developed, which requires large appliance-level data sets for the initial off-line training phase (e.g., Kolter et al., 2010; Elhamifar and Sastry, 2015; Singh et al., 2016). Despite a certain level of intrusiveness is unavoidable to ensure accuracy in the subsequent stages of data disaggregation, the challenge is to keep it at a minimum. This challenge has been motivating the recent emergence of a second class of unsupervised algorithms, which generally avoid collecting appliance-level data (see Bonfigli et al., 2015, and references therein).

Second, the definition of consistent accuracy metrics against which NILM algorithms can be evaluated and compared is another domain challenge. According to Butner et al. (2013), Barker et al. (2014) and Batra et al. (2014), no consistent conventions and standards are currently in place for measuring the accuracy of NILM technologies. Many algorithms tend to focus only on accurately detecting the *on/off* status of each appliance (e.g., Giri and Bergés, 2015; Makonin and Popowich, 2015; Bernard and Marx, 2016) and their accuracy is hence evaluated using metrics accounting for *on/off* detection, such as the F-score (Batra et al., 2014). Only few studies also consider the accuracy in reproducing the consumption patterns of single end-uses in time, which is evaluated either by visual inspection or by means of specific quantitative metrics (Kolter and Johnson, 2011; Gabaldón et al., 2014; Kelly and Knottenbelt, 2015; Bonfigli et al., 2015; Amenta and Tina, 2015; Mueller and Kimball, 2016; Piga et al., 2016). While limiting the extent of NILM algorithms to only the detection of *on/off* events allows retrieving information on appliances time and frequencies of use, a correct reproduction of end-use patterns would support water utilities and demand management with more exhaustive information regarding consumers' behavior and energy usage efficiency. Accurate estimates of appliances power consumption patterns enables a better identification of peak-hours, a more accurate quantification of the power load contributed by each appliance during peak and off-peak hours, as well as assessments on the efficiency levels of different appliances. These are key information to understand consumers' behaviour and, ultimately, design personalized demand management strategies targeted at improving power consumption efficiency and reducing costs, for instance through demand peak-shifting and retrofitting of low-efficiency devices. Finally, a third challenge to energy disaggregation algorithms consists in the number of simultaneously operating appliances that can be identified by NILM algorithms (Froehlich et al., 2011; Butner et al., 2013; Barker et al., 2014). This is a double challenge because an increasing number of simultaneously operating appliances not only raises the variety of appliance-specific consumption pat-

terns to be identified, but also increases the combinations of overlapping uses, and, consequently, signal distortion (Liang et al., 2010).

In this chapter, we address these three challenges by contributing two novel algorithms for NILM. The first algorithm, later on called SOD, exploits sparse optimization to perform end-use disaggregation. It formulates the problem of end-use disaggregation as a least-square error minimization problem, with an additional (convex) penalty term aiming at enforcing the disaggregate signals to be piecewise constant over the time. Moreover, it is based on the assumption that the unknown appliance power/water consumption profiles are piecewise constant over time and it exploits the information on the time-of-day probability in which a specific appliance might be used. The second one, called HSID, combines FHMMs and ISDTW to accurately characterize end-use trajectories for a number of simultaneously operating appliances and reduce the intrusiveness of the off-line training. More precisely, the FHMMs module of the algorithm initially disaggregates the total power/water consumption signal into 2-state single-appliance piece-wise constant trajectories. Thus, FHMMs provide a rough approximation of the end-use trajectories. ISDTW is then applied, in order to reshape them according to the typical power/water consumption pattern of each specific end-use, and include the intrinsic variability of the latter in terms of power/water range and appliance usage duration. After being processed through ISDTW, the estimated end-use trajectories describe more accurately and realistically the power/water consumption time series of each appliance. We present here a supervised and a semi-supervised versions of the algorithm, developed in order to deal both with applications involving intrusive measurements at single-appliance level, as well as non-intrusive ones. The two versions of HSID are independent and differentiate with respect to the information needed for algorithm training: the supervised version of HSID requires appliance-level load measurements, while the semi-supervised version exploits aggregate measurements from the smart meter to retrieve appliance-level information.

This chapter is organized as follows. We formalize the disaggregation problem in Section 3.2 and describe the two algorithms in Section 3.3. In Section 3.4, we present multiple tests of the two algorithms on real-world power consumption data. In particular, we first assess their performance through diverse set of metrics and against a state-of-the-art benchmark method. We then test the sensitivity of HSID with respect to the level of noise in the metered consumption as well as the number of metered appliances. Moreover, we demonstrate the usability of HSID in semi-supervised applications, without the need of gathering a training dataset at the appliance level, and test the effect of lowering data resolution on disaggregation accuracy. Finally, in Section 3.5 we assess the portability of HSID to the disaggregation of synthetic water data and then we

conclude the chapter with a general discussion on all the results obtained.

3.2 Problem formulation and related work

NILM end-use disaggregation algorithms estimate the power consumption of each appliance contributing to the total consumption of a household as measured by a single-point smart meter at sub-daily frequency, namely either low-frequency (e.g., 10 seconds or 1 minute) or high-frequency (e.g., hundreds of Hz).² This problem can be classified as a *blind identification* problem (Abed-Meraim et al., 1997) where, given the observed output of the whole system (i.e., the household total power consumption), the unobserved sub-states (i.e., the power consumption of each appliance) should be estimated. More formally, we can write the total power consumption of a house at time step t as:

$$\bar{Y}_t = \sum_{i=1}^N y_t^i + e_t \quad (3.1)$$

where \bar{Y}_t is the total, observed power consumption at each time step t , y_t^i the consumption of appliance i at time step t , N the total number of appliances, and e_t is the measurement noise.

The consumption of the i -th appliance is written, for each time step, as:

$$\begin{aligned} y_t^i &= \mathbf{B}^i \mathbf{x}_t^{i,T} + \epsilon_t^i & y_t^i &\in \mathbb{R}^+ \\ \mathbf{B}^i &= [b^{i,1}, b^{i,2}, \dots, b^{i,M}] & i &= 1, \dots, N \\ \mathbf{x}_t^i &= [x_t^{i,1}, x_t^{i,2}, \dots, x_t^{i,M}] & i &= 1, \dots, N \end{aligned} \quad (3.2)$$

where:

- \mathbf{B}^i is a vector containing the power consumption basis $b^{i,j}$ for each appliance i , i.e., the power consumption related to each operating state j (e.g., *on/off*) of the appliance. M is the number of potential power states (in this formulation it is assumed that all the appliances have the same number of states);
- \mathbf{x}_t^i represents the activation vector for the states of appliance i , at time t . It is a binary-valued vector indicating which power levels of vector \mathbf{B}^i are operating for the i -th appliance at each time step, therefore $x_t^{i,j} \in \{0, 1\} \forall i, j, t$. Also, each appliance can only operate in one state at a time, thus the following constraint holds: $\sum_{j=1}^M x_t^{i,j} = 1 \quad \forall i, t$;

²We remind that in the theoretical sections of this chapter we refer to power consumption disaggregation, but the same concepts stand for water consumption disaggregation, where not specified.

- ϵ_t^i is the noise affecting the consumption of appliance i at time step t . It may be due to intrinsic characteristics of appliances or mutual interference among appliances on the same network.

Based on the elements introduced so far, we can formulate the general disaggregation problem as a minimization problem, searching for the consumption trajectories of each appliance that minimize the error between estimated and real power trajectories, for each time step of the considered time horizon H :

$$\begin{aligned}
 [\mathbf{B}^*, \mathbf{x}_t^*] &= \arg \min_{\mathbf{B}, \mathbf{x}_t} \left[\sum_{t=1}^H (\bar{Y}_t - \hat{Y}_t)^2 \right] \\
 &= \arg \min_{\mathbf{B}, \mathbf{x}_t} \sum_{t=1}^H \left[\sum_{i=1}^N (\bar{y}_t^i - \hat{y}_t^i)^2 \right]
 \end{aligned} \tag{3.3}$$

The above formulation defines the general problem of the NILM methods classified in Zoha et al. (2012) as *optimization methods*. In principle, the solution to Problem 3.3 may be computed by standard least-square techniques. However, the problem is overparametrized (i.e., the number of parameters to be estimated is higher than the number of available measures) and, in practice, alternative *optimization methods* such as genetic algorithms (Baranski and Voss, 2004), integer optimization (Suzuki et al., 2008) or sparse optimization (Kolter et al., 2010; Elhamifar and Sastry, 2015; Piga et al., 2016) are used to search for the best match between a combination of appliances sampled from a known database and the vector of total measured consumption. Despite showing good accuracy in cases with a limited number of appliances combinations, the computational complexity and the lack of inclusion of the temporal continuity of power signals constitute two major drawbacks for such techniques. SOD, the first NILM algorithm that we propose in this chapter, can be classified as an optimization methods, as it exploits sparse optimization techniques to perform disaggregation. The issues of temporal structure, continuity, and state transition are tackled by the so called *pattern recognition* methods (Zoha et al., 2012). The algorithms belonging to this class do not approach Problem 3.3 as an independent problem for each time step. In contrast, they include information about the temporal structure of power signals and search for the sequence of appliances states that is optimal with respect to those features that show temporal continuity (e.g., state transition probabilities). Techniques such as Artificial Neural Networks (ANNs) (Srinivasan et al., 2006) or HMMs (Kolter and Jaakkola, 2012) have been tested in this context, successfully showing the value of temporal information in learning the consumption patterns of appliances. In particular, several HMM-based load disaggregation algorithms have been widely used and discussed in NILM literature (Mauch et al., 2016), showing the

3. Water and energy end-use disaggregation

potential of achieving disaggregation accuracies higher than 70% and up to 99% on relevant appliances, when they are trained on data from the same household used for testing. Given these promising results, HMM-family methods (Kolter and Johnson, 2011; Batra et al., 2014) are often adopted as benchmarks for algorithm testing and comparison. Another family of recently proposed pattern recognition methods relies on DTW, a well-established pattern matching technique used in the field of speech recognition, which allows comparisons and matching between traces of different length (Sakoe and Chiba, 1978). In the field of electricity smart grids, DTW has been mainly used for electric profiles clustering (Gullo et al., 2009).

All these algorithms rely on a preliminary supervised learning phase, which requires to intrusively collect key information, such as the number of appliances N or the typical end-use power consumption pattern of each appliance (the so-called *signatures*), for the calibration of the power consumption bases \mathbf{B} and other method-specific parameters. The supervised version of the proposed HSID algorithm can be hence classified as a supervised pattern matching method, which combines FHMMs and ISDTW and requires the knowledge of N and of the appliances' signatures. In particular, N can be either retrieved intrusively through door-to-door appliance surveys or can be reported directly by householders. Especially the latter case, or when the number of electric appliances is large, the reported values of N may be affected by errors, that will be then propagated along the disaggregation process, if no cross-checks are considered. However, the intrusiveness of collecting appliance-level data for training the disaggregation algorithms often hampers the wide usability of such tools in real world applications. A non-intrusive alternative is represented by completely unsupervised algorithms (see Bonfigli et al., 2015, and references therein), which do not require any *a-priori* information about N and \mathbf{B} . These methods have been benchmarked against well known datasets and supervised methods (see Kim et al., 2011; Goncalves et al., 2011; Shao et al., 2012; Liao et al., 2014; Parson et al., 2014; Zhao et al., 2015; Pöchacker et al., 2016; Liu et al., 2016; Bernard and Marx, 2016, for instance). Yet, they are subject of ongoing research in the field of electrical power disaggregation, and, given the recent development of most of them, comprehensive studies cross-comparing the performance of promising unsupervised methods against common datasets and performance metrics are in progress (Bonfigli et al., 2015). An alternative approach for reducing the intrusiveness of supervised methods is to extract appliance-level signatures from the aggregate-level smart metered trace (Parson et al., 2014). In practice, the availability of consumption diaries (Zhao et al., 2016), where the time-of-use of each appliance is recorded for part of the monitoring period, would support this approach. Recently, Elafoudi et al. (2014) proposed a NILM algorithm based on DTW and customers' daily diaries,

which is demonstrated to attain promising results in the end-uses disaggregation, with an overall disaggregation accuracy in terms of F-score higher than 85%. Indeed, power use diaries information would allow retrieving appliances signatures from portions of the total consumption trace when only one appliance is operating (i.e., no other appliances are simultaneously in operation). We developed a second, semi-supervised version of the proposed HSID algorithm upon this idea, and show that the algorithm is portable in almost non-intrusive cases, where only the knowledge of the number of appliances N is needed, thus removing the initial intrusiveness and costs associated with the metering of fixture-level data. Both SOD and HSID algorithms are described in the next section.

3.3 Methods

3.3.1 Sparse optimization-based algorithm

3.3.1.1 Theoretical concepts

In this section, the sparse optimization based NILM algorithm SOD is presented. The following conditions are assumed to hold for its development:

Assumption 1: a training data set D_{T_t} is available. The training set consists of the observations of the power signatures of each appliance available in the house. An intrusive period is needed to construct D_{T_t} . During this period, the patterns of electricity demand of each appliance are observed, and information on time-of-day probability characterizing the usage of each appliance can be also gathered.

Assumption 2: a roughly knowledge of the power demand of each appliance at each operating mode (i.e., the terms $b_{i,j}$ in Equation 3.2) is supposed to be available. For instance, the terms $B_i^{(j)}$ can be evaluated from D_{T_t} through k-means clustering Likas et al. (2003) or through a simple visual inspection.

Assumption 3: the energy consumption profiles of each appliance are piecewise constant over time.

Standard Least Squares In order to estimate the power demand y_t^i of each appliance based on the aggregate power consumption observations D_{T_t} , the time-varying parameters $\chi_t^{i,1}, \chi_t^{i,2}, \dots, \chi_t^{i,M}$ (with $i = 1, \dots, N$ and $t = 1, \dots, T$)

3. Water and energy end-use disaggregation

can, in principle, be computed solving the standard least-square problem:

$$\min_{\{\mathbf{x}_t^{i,j}\}} \sum_{t=1}^T \left(\bar{Y}_t - \sum_{i=1}^N \hat{y}_t^i \right)^2 \quad (3.4)$$

$$\begin{aligned} i &= 1, \dots, N \\ j &= 1, \dots, M \\ t &= 1, \dots, T \end{aligned}$$

with

$$\hat{y}_t^i = y_t^i - \epsilon_t^i = \mathbf{B}^i \mathbf{x}_t^{i,\top}$$

However, Problem (3.4) is overparameterized, since it involves more parameters than measurements. As a consequence, overfitting occurs, and thus no generalization property is guaranteed. One possible solution to overcome this problem is to introduce regularization terms in (3.4) to:

- enforce each appliance to operate at a single mode at each time instant;
- according to **Assumption 3**, enforce the energy consumption profiles \hat{y}_t^i to be piecewise constant signals over time.

Adding regularization In order to exploit the information that: (i) the parameters $x_t^{i,1}, x_t^{i,2}, \dots, x_t^{i,M}$ can be either 0 or 1 and (ii) each appliance can only operate at a single mode at each time instant, the following regularized problem can be solved instead of (3.4):

$$\min_{\{\mathbf{x}_t^{i,j}\}} \sum_{t=1}^T \left(\bar{Y}_t - \sum_{i=1}^N \hat{y}_t^i \right)^2 + \lambda_1 \sum_{i=1}^N \sum_{t=1}^T \left\| \begin{bmatrix} x_t^{i,1} \\ \vdots \\ x_t^{i,M} \end{bmatrix} \right\|_0 \quad (3.5)$$

$$\begin{aligned} i &= 1, \dots, N \\ j &= 1, \dots, M \\ t &= 1, \dots, T \end{aligned}$$

s.t.

$$\sum_{j=1}^M x_t^{i,j} = 1 \quad \forall i, t,$$

$$x_t^{i,j} \in \{0, 1\} \quad \forall i, j, t$$

where $\|\cdot\|_0$ denotes the cardinality of a vector (i.e., number of its nonzero components). Note that, on one hand, the second term in the objective function of Problem (3.5) aims at enforcing sparsity in the vector \mathbf{x}_t^i . On the other hand, the vector \mathbf{x}_t^i is guaranteed to have at least one element different than

zero, because of the equality constraint appearing in Problem (3.5). The hyperparameter $\lambda_1 > 0$ is tuned by the user (for instance, through cross-validation) for balancing the tradeoff between minimizing the fitting error (by decreasing the value of λ_1) and maximising sparsity of the parameter vector \mathbf{x}_t^i (by increasing the value of λ_1).

Note that, because of the $\|\cdot\|_0$ operator, Problem (3.5) is not convex. According to the Lasso (Tibshirani, 1996; Zou, 2006), an approximate solution of Problem (3.5) can be obtained by replacing the cardinality of a vector (i.e., the operator $\|\cdot\|_0$) with its ℓ_1 norm. Furthermore, in order to improve the accuracy of the final estimate, the parameters \mathbf{x}_t^i can be scaled by nonnegative weights \mathbf{w}_t^i . This leads to the following convex approximation of Problem (3.5):

$$\begin{aligned} \min_{\substack{\{\mathbf{x}_t^{i,j}\} \\ i=1,\dots,N \\ j=1,\dots,M \\ t=1,\dots,T}} \quad & \sum_{t=1}^T \left(\bar{Y}_t - \sum_{i=1}^N \hat{y}_t^i \right)^2 + \lambda_1 \sum_{i=1}^N \sum_{t=1}^T \left\| \begin{bmatrix} w_t^{i,1} \\ \vdots \\ w_t^{i,M} \end{bmatrix} \odot \begin{bmatrix} x_t^{i,1} \\ \vdots \\ x_t^{i,M} \end{bmatrix} \right\|_1 \end{aligned} \quad (3.6)$$

s.t.

$$\begin{aligned} \sum_{j=1}^M x_t^{i,j} &= 1 \quad \forall i, t, \\ x_t^{i,j} &\in \{0, 1\} \quad \forall i, j, t \end{aligned}$$

where \odot denotes the element-wise multiplication. The choice of the weights \mathbf{w}_t^i is discussed in Section 3.3.1.2. Note that the ℓ_1 -norm regulation promotes sparsity of the vector \mathbf{x}_t^i . In fact, in the ideal case, only one component of that vector should be nonzero (i.e., the i -th appliance operates at a single mode at each time instant). The reader is referred to the works Tropp (2004); Fuchs (2004); Donoho (2006); Vincent and Novara (2013) for a detailed analysis on the properties of ℓ_1 -regularization in sparse estimation problems.

Adding regularization to enforce piece-wise constant signals power demand profiles In order to further improve the accuracy of the estimate given by (3.6), we might exploit the additional information that the power demand signatures of the electric appliances are piecewise constant over time (**Assumption 3**). In order to enforce the power signals to be piecewise constant, a new regularization term aiming at penalizing the variation of the time-varying

3. Water and energy end-use disaggregation

coefficients \mathbf{x}_t^i is added to Problem (3.6), i.e.,

$$\begin{aligned} \min_{\substack{\{\mathbf{x}_t^{i,j}\} \\ i = 1, \dots, N \\ j = 1, \dots, M \\ t = 1, \dots, T}} \quad & \sum_{t=1}^T \left(\bar{Y}_t - \sum_{i=1}^N \hat{y}_t^i \right)^2 + \lambda_1 \sum_{i=1}^N \sum_{t=1}^T \left\| \begin{bmatrix} w_t^{i,1} \\ \vdots \\ w_t^{i,M} \end{bmatrix} \odot \begin{bmatrix} x_t^{i,1} \\ \vdots \\ x_t^{i,M} \end{bmatrix} \right\|_1 \\ & + \lambda_2 \sum_{i=1}^N \sum_{t=1}^T \left\| k_i \begin{bmatrix} x_t^{i,1} - x_{t-1}^{i,1} \\ \vdots \\ x_t^{i,M} - x_{t-1}^{i,M} \end{bmatrix} \right\|_\infty \end{aligned} \quad (3.7)$$

s.t.

$$\begin{aligned} \sum_{j=1}^M x_t^{i,j} &= 1 \quad \forall i, t, \\ x_t^{i,j} &\in \{0, 1\} \quad \forall i, j, t \end{aligned}$$

with λ_2 being a tuning parameter playing a role similar to λ_1 . The terms k_i (with $i = 1, \dots, N$) are *a-priori* specified nonnegative weights which can be chosen through the method described in Section 3.3.1.2. It is worth remarking that:

- penalizing the norm of the difference between two consecutive parameters $x_t^{i,j} - x_{t-1}^{i,j}$ is commonly referred in the literature as *Fused Lasso* (Tibshirani et al., 2005) and it is used to promote sparsity in the discrete-time derivative of the signal $x_t^{i,j}$ (thus enforcing the signal $x_t^{i,j}$ to be piecewise constant over time).
- the last term of Equation (3.7) is a *group (fused) Lasso* penalty (Yuan and Lin, 2006; Zhao et al., 2009; Vogt and Roth, 2010), penalizing the mixed $\ell_{1,\infty}$ -norm (i.e., sum of the infinity norms) of the groups

$$\begin{bmatrix} x_t^{i,1} - x_{t-1}^{i,1} \\ \vdots \\ x_t^{i,M} - x_{t-1}^{i,M} \end{bmatrix},$$

with $i = 1, \dots, N$ and $t = 2, \dots, T$. The infinity norm is considered for this last term so that, at the solution, the vector

$$\begin{bmatrix} x_t^{i,1} - x_{t-1}^{i,1} \\ \vdots \\ x_t^{i,M} - x_{t-1}^{i,M} \end{bmatrix},$$

is enforced to be either identically zero or full. In fact, if one of the parameters in \mathbf{x}_t^i changes from time $t - 1$ to t , a variation of the other parameters does not change the cost function. Specifically, only the largest time variation among the elements of the vector \mathbf{x}_t^i affects the objective

function. Following the same rationale, the ℓ_2 -norm can be alternatively used instead of the ℓ_∞ -norm. The choice of the norm of the group is a problem at hand, mainly related to the numerical algorithms used to solve the formulated group Lasso problem.

Exploiting the information that every appliance cannot change state simultaneously If the sampling interval $\Delta_t = t - (t - 1)$ is small enough, it is also reasonable assuming that at most one appliance can change operating mode at each time instant. If this assumption holds, this additional information can be exploited by adding the following convex constraints to Problem (3.7):

$$\sum_{i=1}^N \left\| \begin{bmatrix} \chi_t^{i,1} - \chi_{t-1}^{i,1} \\ \vdots \\ \chi_t^{i,M} - \chi_{t-1}^{i,M} \end{bmatrix} \right\|_\infty \leq 1, \quad t = 2, \dots, T \quad (3.8)$$

3.3.1.2 Practical implementation

Some suggestions for a practical implementation of the proposed disaggregation algorithm, including the choice of the weighting parameters \mathbf{w}_t^i and k_i from the training set D_{T_t} , are given in this section.

\mathbf{w}_t^i weights setting The main idea behind the choice of the weights $w_t^{i,1}, \dots, w_t^{i,M}$ is the following: if the i -appliance is more likely to operate at mode j at time t , then the parameter $\chi_t^{i,j}$ is more likely to be equal to 1, while the other parameters $\chi_t^{i,g}$ (with $g \neq j$) are more likely to be equal to zero. In terms of the optimization problem (3.7), the parameters $\chi_t^{i,g}$ (with $g \neq j$) should be more penalized than $\chi_t^{i,j}$, or equivalently, $w_t^{i,g}$ (with $g \neq j$) should be higher than $w_t^{i,j}$. The information on time-of-day probability of the usage of each appliance can be inferred from the training data set D_{T_t} . Specifically, for given i and t , the weights $w_t^{i,1}, \dots, w_t^{i,M}$ can be chosen as follows:

1. given D_{T_t} , for each time sample t compute $q_t^{i,j}$ as the number of times the i -th appliance is classified to be at mode j at the time samples $t + k24h$, with $k \in \mathbb{Z}$.
2. if $q_t^{i,j} \neq 0$, the weight $w_t^{i,j}$ is then given by: $q_t^{i,j} = \frac{1}{q_t^{i,j}}$. Otherwise set the weight $w_t^{i,j}$ to a large number.

Note that the parameter $q_t^{i,j}$ might also be computed considering not only the observations at $t, t - 24h, t + 24h, t - 48h, t + 48h, \dots$, but also the observations (possibly weighted) within given time intervals $[t + k24h - \Delta, t + k24h + \Delta]$.

3. Water and energy end-use disaggregation

k_i weights setting The weights k_i can be chosen as follows: if the i -th appliance rarely changes its operating mode over time, than the time variation of the parameters $\chi_t^{i,j}$ should be more penalized w.r.t. the time variation of the parameters characterizing another appliance which frequently changes its operating mode. The weight k_i can be then inversely proportional to the number of mode changes observed in the training dataset for the i -th appliance, and scaled by the length of the training dataset.

Reducing the computational complexity As the number of optimization variables in Problem (3.7) grows linearly with the length T of the signal \bar{Y}_t to be disaggregated, the applicability of the proposed approach might be limited to small/medium values of T . In order to overcome this problem, a sub-optimal solution of Problem (3.7), can be simply computed by splitting the dataset D_T into H disjoint subsets \mathcal{D}^h of length T_h (with $h = 1, \dots, H$) such that $D_T = \bigcup_{h=1}^H \mathcal{D}^h$. Problem (3.7) is then solved for each subset \mathcal{D}^h .

The computational complexity of the algorithm can be further reduced as follows. If at time t the i -th appliance is guaranteed not to operate at the j -th mode, then the parameter $\chi_t^{i,j}$ can be set to zero, thus reducing the number of decision variables for Problem (3.7). Such an information can be simply obtained by analyzing the observed aggregate power consumption \bar{Y}_t . In fact, if $\bar{Y}_t \ll b^{i,j}$ (i.e., the observed aggregate power consumption at time t is largely lower than the power consumption of the i -th appliance when operating at mode j), then $\chi_t^{i,j}$ can be directly set to zero.

3.3.2 Hybrid Signature-based Iterative Disaggregation algorithm

The development of the second algorithm for end-use disaggregation that we propose here, Hybrid Signature-based Iterative Disaggregation, is based on the following two assumptions:

Assumption 1: each electrical device contributing to the total household consumption can be recognized from its specific consumption pattern, i.e., each fixture has a typical *signature* (Ruzzelli et al., 2010; Dong et al., 2013), such as the example in Figure 3.1.

Assumption 2: the consumption level time series of each appliance can be modeled as a Markovian state sequence, and can be represented with a limited number of states (e.g., state 1: fixture on/operating; state 2: fixture off/not operating). The two assumptions might not hold for appliances with extremely noisy and scattered behavior or unrealistic situations characterized by random shifts in single- appliance operating range and operating status transitions.

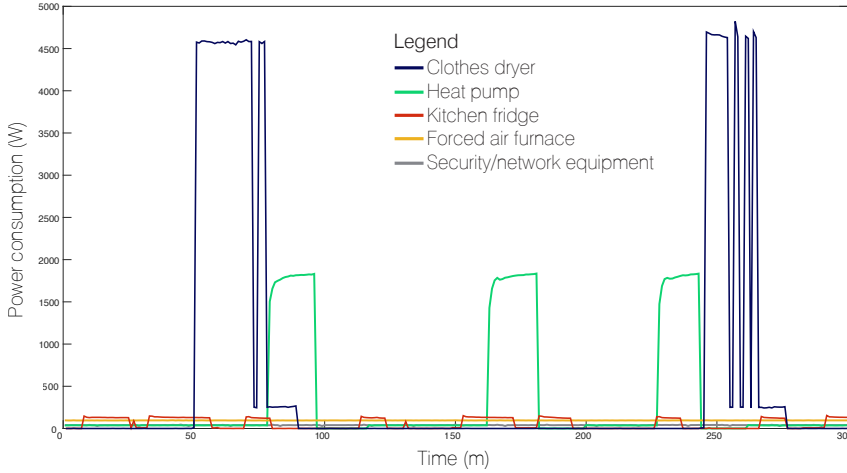


Figure 3.1: Typical power load signatures for five indoor appliances.

Operationally, the workflow of the algorithm is composed of the following three steps (Figure 3.2):

- A. appliances signatures identification;
- B. disaggregation of household power load through 2-state Factorial Hidden Markov Models;
- C. end-use trace patterns correction through Iterative Subsequence Dynamic Time Warping.

Two versions of the HSID algorithm are presented in this chapter: the first one, *supervised*, requires initial intrusive appliance-level load measurements for algorithm training, while the second, *semi-supervised*, retrieves appliance-level information directly from the aggregate measurements provided by the smart meter.

We provide details on each step of the HSID algorithm for both settings in the next paragraphs.

3.3.2.1 Supervised HSID algorithm

A. Appliances signatures identification This first step of the supervised version of HSID aims at creating a database containing the signature s^i of each appliance i , ($i = 1, \dots, N$) contributing to the total measured consumption at the household level. In order to gather all the needed signatures s^i , the algorithm requires a training dataset D_T , consisting of power consumption observations retrieved intrusively for each appliance/fixture during a short training period. The length of the training period T should be kept as short as possible,

3. Water and energy end-use disaggregation

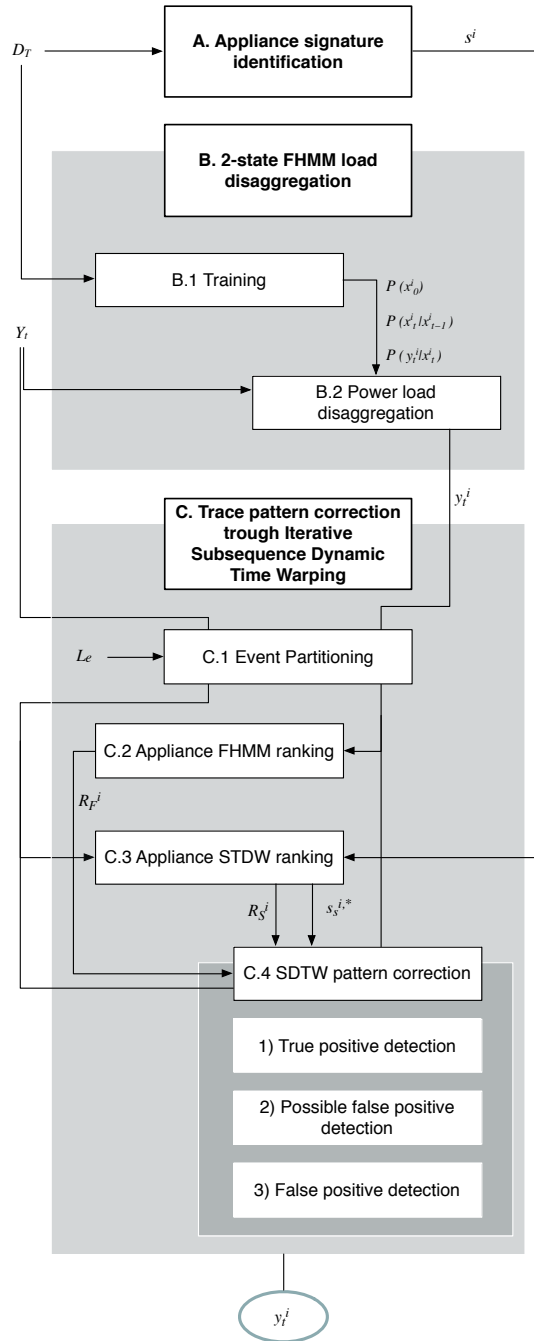


Figure 3.2: Flowchart of the HSID algorithm. Boxes enumeration is consistent with the one adopted in Section 3.3.2.

in order to reduce intrusiveness and costs. In this study we considered a 2-week long training period, in order to take into account possible consumption differences between week and weekend days and gather a meaningful sample of power consumption events. We implemented the signature identification in two steps. In the first, the consumption trajectory $y^{i,T} = \{\bar{y}\}_{t=0,\dots,T}^i$ of each appliance during the training period is retrieved from D_T . As a second step, each signature s^i is obtained by removing the trace noise (Dong et al., 2013) from the lowest values of $y^{i,T}$. Those values are identified through a 3-clusters K-means algorithm (MacQueen et al., 1967) and are set all equal to the centroid of the lowest cluster, so that no background noise affects the quality of signatures.

B. 2-state FHMM load disaggregation The purpose of this step is to perform the power load disaggregation. The total power consumption trace of each household is partitioned into simplified two-state consumption trajectories for each fixture, in order to identify their *on/off* operating states. We used FHMMs for this purpose. FHMMs (Ghahramani and Jordan, 1997) are a well established technique in machine learning and have already been applied in the field of power load disaggregation (Batra et al., 2014). FHMMs allow the identification of the most probable sequence of states of a Markovian process when the considered system is composed of different sub-components, and the state of the whole system (i.e., the only measured element) is a combination of the hidden states of each sub-component. In FHMMs, each of the N appliances is characterized by a finite number of hidden states (i.e., not observed), the latter described by a prior probability distribution, and a matrix with probabilities of transition between couples of states. Each appliance can thus be modeled with a Hidden Markov Model, and then single-appliance HMMs are combined in a FHMM, considering that there is a specific probabilistic relation between the observation and combinations of hidden states (i.e., *emission probability*). In this work, we use the NILM toolkit proposed by Batra et al. (2014), which implements a FHMM with Gaussian emission probabilities and exact inference (Ghahramani and Jordan, 1997).³ FHMM performs power load disaggregation according to the following two-step procedure:

B.1 Training. The training phase of FHMM considers the same training dataset D_T used for the initial signatures extraction phase, and consists in the calibration of the three main elements of HMMs: (i) the *marginal initial probability distribution* $P(\mathbf{x}_0^i)$ for each appliance, i.e., the probability of occurrence of each state in the initial time step, (ii) the *transition probability distribution* $P(\mathbf{x}_t^i | \mathbf{x}_{t-1}^i)$ for the states of each appliance, i.e., the transition probability among the different operating modes of each appliance between two sequential time steps, and (iii) the *emission probability* distribu-

³The HMM module of Python scikit-learn machine learning library is exploited for FHMM resolution.

3. Water and energy end-use disaggregation

tion $P(y_t^i | x_t^i)$ for the states of each appliance, i.e., the probability of observing a particular output of the system depending on its operating state.

B.2 Power load disaggregation. Once the above three probability distributions are calibrated, FHMM performs a disaggregation of the total consumption data over the validation time horizon H . In short, FHMM solves the following instance of Problem 3.3:

$$[P^*(x_t^i | x_{t-1}^i), P^*(y_t^i | x_t^i)] = \arg \min_{P^*(x_t^i | x_{t-1}^i), P^*(y_t^i | x_t^i)} \left[\sum_{t=1}^H (Y_t - \hat{Y}_t)^2 \right] \quad (3.9)$$

and then it uses the Viterbi algorithm (Viterbi, 1967) to identify the most probable sequence of (hidden) states associated with the measured output.

For computational reasons, our model assumes that the number of states each appliance has in the FHMM is equal to 2 (for more details, see Section 3.4.2.1). As a consequence, the consumption trajectories estimated for each appliance by FHMM assume the shape of piecewise constant lines, i.e., only the on/off operating states are detected, while an accurate reproduction of power consumption patterns is missing at this stage. In addition, this two-state outcome is not acceptable to accurately reproduce the consumption patterns of multi-state appliances, such as washing machines, or *Continuously Variable Devices* (Zoha et al., 2012), whose behaviour cannot be captured by a two-state sequence. The challenge of retrieving the variety of consumption patterns for such appliances and avoiding estimation error propagation due to oversimplified trajectories is tackled by the last component of our algorithm.

C. Trace patterns correction through Iterative Subsequence Dynamic Time Warping In this phase, we iteratively use Subsequence Dynamic Time Warping (SDTW) (Sakoe and Chiba, 1978; Müller, 2007) in HSID to integrate the information on the consumption patterns variety given by the signatures extracted at the beginning of the procedure (Step A), and to correct the 2-state trajectories produced as output in the FHMM step (Step B). We integrated SDTW pattern-matching technique in our algorithm according to the following procedure:

C.1 Event partitioning. The total consumption trajectory and the single appliances trajectories estimated by FHMM are split into time windows of equal length, henceforth called *events*. The length of the event L_e is tuned as an average of the durations of the pulses in the total consumption trajectory, in order to be consistent with the typical power usage durations in the dataset.

C.2 Appliance FHMM ranking. For each event, the appliances are ranked in decreasing order according to the values of the 90-th percentile of their FHMM trajectories within the event. Each appliance is labelled with an

ordinal value R_F^i . This rank gives an idea of the contribution each appliance brings to the total event, so that the trajectories of the highest ranked appliances, i.e., the ones with a larger power contribution in the considered event, can be corrected first, as they have a largest impact on total consumption and “hide” the trajectories of less contributing appliances.

C.3 Appliance SDTW ranking. For each event, appliances are assigned a second ranking R_S^i . This is computed by evaluating the similarity between the observed total power consumption trace of the considered event and the signature s^i of each appliance. The similarity is given by the distance among the two trajectories as evaluated by DTW: the larger the distance, the lower the similarity of the two trajectories. Therefore, highest ranked signatures are the ones closest to the power trace of the event. DTW is applied as a SDTW (Müller, 2007), because the length of events is usually much shorter than the total length of the signature. Indeed, signatures are extracted from appliance-level traces long as much as the training period, as explained at point A of this procedure, while consumption events usually last a few minutes. Therefore, signatures can potentially contain more than one consumption event: this means that the signature is not entirely compared to the total power consumption trace of the event, but it is scanned in order to find the best event matching sub-sequence $s_s^{i,*}$.

C.4 SDTW pattern correction. In this very last phase of the algorithm, the 2-state power load trajectories estimated by FHMM are corrected keeping into account the information given by the signatures and the ranking vectors R_F^i and R_S^i . SDTW pattern matching is iteratively applied according to the following alternative three-case heuristics:

1. *True positive detection.* If $R_F^1 = R_S^1$, i.e., an appliance is ranked first both by FHMM and DTW rankings, the *on/off* operating state detected by FHMM is assumed to be correct. Given that, the estimated consumption pattern for that appliance can be refined by replacing the piecewise constant Markov state with the best matching portion of signature $s_s^{1,*}$ found at point C.3.
2. *Possible false positive detection.* If $R_F^1 \neq R_S^1$, meaning the first ranked appliance by FHMM is not the one with the most similar signature to the total load power, FHMM may have generated a *false on* event. In this uncertain situation, we assume that the state activation is inertial (Kolter et al., 2010), i.e., if an appliance was *on* at time $t - 1$ is likely to be *on* again at time t and viceversa. This assumption penalizes unrealistic frequent state transitions (e.g., continuously turning *on* and *off* an appliance in a short time interval) and is implemented as follows: if the fixture ranked as R_F^1 had a larger normalized contribution to the total power consumption over the previous k time steps than the one

3. Water and energy end-use disaggregation

of the appliance ranked as R_S^1 , it is kept as active and corrected with its signature $s_S^{i,*}$ (see step 1). Elsewhere, it is switched off according to the next case.

3. *False positive detection.* If none of the previous cases is met, it is assumed that the *on* state generated by FHMM for the considered appliance is a false positive, therefore the appliance is switched off and its consumption trajectory is set to its lowest Markov state.

After the first ranked appliance is corrected, the residual total power consumption is updated and step C.4 is repeated recursively for the remaining appliances. Thus, the procedure is iterated in order to correct the signal of all the simultaneously operating appliances, without requiring only one appliance operating at each time step.

It is important mentioning that an exception holds for corrections 1 and 2: signature correction is not implemented when its implementation would introduce noise on the signal estimated by FHMM, i.e., when the total consumption signal is more similar to the estimated 2-state appliance signal than to any signature s^i . A further exception holds for appliances with a training trajectory confined in a very narrow power interval (i.e., lower or equal to the lowest positive Markov state in the state space of the problem): those appliances are corrected prior than the other ones and set constantly equal to the average of their signature, because this latter is noisy but varies within a narrow interval. Finally, it is worth noticing that no false negative cases are explicitly considered, as they are automatically solved by difference, given that each iteration considers the updated residual of the total power load for ranking and similarity with signatures.

Overall, we expect HSID to benefit from ISDTW in order to produce accurate disaggregated end-use trajectories. In a previous work by Elafoudi et al. (2014), DTW was successfully adopted to classify energy consumption events by matching with labeled templates stored in a reference library. This was demonstrated to achieve high performance in appliance usage detection. In HSID, we extend the use of DTW so that it exploits signature information to shape and correct FHMM-estimate trajectories and increase the accuracy in end-use trajectory estimation, consequently providing better estimate of the actual amount of energy used by each appliance, as well as reducing event detection errors. The ISDTW we include in HSID can process long signatures comprising multiple usage events, in order to find the best matching portion of each signature (subsequence) and, afterwards, exploit the latter to correct and refine the disaggregated end-use trajectories. As HSID iteratively repeats this process for each event on the residual power load until all the load of the event has been disaggregated among different appliances, we adopt ISDTW in a decomposition,

rather than classification, mode.

3.3.2.2 Semi-supervised HSID algorithm

Following the limitations posed by the intrusiveness of the supervised learning phase, in this paper we propose also a semi-supervised (i.e., appliance -level measure free) version of the HSID algorithm. This second version of our algorithm does not require appliance-level ground-truth training data and manipulates only the total energy consumption metered at the household level. In this semi-supervised scenario, we assume that a single-event signature, i.e., the signature of each appliance for a single event, can be retrieved from the total power consumption pattern upon knowledge of a time window in which the considered appliance is working without other appliances interference. This is equivalent to a situation in which no on-device smart sensors are installed, thus avoiding the intrusiveness from the point of view of measurements and hardware, but energy activities diaries (Desmedt et al., 2009) filled by energy consumers for a very limited time are available. This situation is realistic, as many energy utilities and multi-utilities worldwide are developing web portals to interact with their customers and provide them with customized services. Energy consumption diaries can be easily included in such portals and users can be allowed to insert information on their consumption (timing and type of device used) as an opt-in, therefore their privacy is safeguarded and intrusiveness avoided. The depicted scenario reflects in the following modifications of HSID with respect to the supervised version described in the previous section: (i) appliances signatures to feed the ISDTW module of HSID are not retrieved from an intrusively gathered appliance-level training dataset as explained in step A of Section 3.3.2, but each signature s^i consists of a single-event signature extracted from the total household power consumption trace, when no other appliances are simultaneously working and (ii) the input dataset D_T for training the FHMM module (see step B of Section 3.3.2) is the union of such signatures $D_T = s^1, s^2, \dots, s^N$.

3.4 Applications to power load end-use disaggregation

We tested and validated the two algorithms presented in the previous paragraphs against sub-daily power consumption data and then extended the usability of HSID onto sub-daily water data. We tested the two algorithms onto power load data before extending their application to water data because the availability of state-of-the-art algorithms for end-use disaggregation, as well as reference datasets, in energy-related research allow for rigorous benchmarking. In particular, the reader will notice that, after testing both algorithms with

3. Water and energy end-use disaggregation

similar experiments on 1-minute resolution power load data, we will perform advanced testing both with electricity and water data only for the HSID algorithm. Our choice is based on two facts: (i) HSID is less demanding in terms of parameters calibration (differently from HSID, SOD requires the calibration of the appliance-specific weights \mathbf{w}_t^i , as well as calibration of the two regularization parameters λ_1 and λ_2), thus requiring less training data than SOD and, (ii) preliminary testing proved that HSID is more computationally efficient.

3.4.1 Power load disaggregation through SOD

3.4.1.1 Experimental settings

Data We tested the SOD algorithm against the AMPDs dataset (Makonin et al., 2013), which contains the power consumption readings of a single house located in the Vancouver region in British Columbia (Canada). Data metered at the end-use level at 1 minute resolution are available for 1-year time, from April 1st 2012 to March 31st 2013. For the sake of analysis and algorithm comparison, we consider only the aggregate power consumption given by the sum of the power consumption readings of the following four electric appliances: clothes dryer; fridge; dishwasher; heat pump. The contribution of the selected appliances is about 45% of the total energy consumption. Furthermore, in order to assess the robustness of the disaggregation algorithm w.r.t. the measurement noise, a fictitious white noise e_t with Gaussian distribution $\mathcal{N}(0, \sigma_e^2)$ and standard deviation $\sigma_e = 4$ W is added to the aggregate power consumption signal y_t . The AMPDs dataset is divided as follows:

- a training set, which consists of the power readings from May 17, 2012 to May 29, 2012. The training set is used to estimate the power demand of each appliance at each operating mode (i.e., the terms in \mathbf{B}^i). Therefore, the sub-metered power consumption trajectories of each appliance are supposed to be available in the training phase. Specifically, the set of power demands of each appliance at each operating mode are chosen through a simple visual inspection of the sub-metered power consumptions in the training dataset. The chosen values are:
 - 1) clothes dryer: [0 260 4700] W,
 - 2) kitchen fridge: [0 128 200] W,
 - 3) dishwasher: [0 120 800] W,
 - 4) heat pump: [39 1900] W.

The training set is also used to estimate the weights w_t^i and k_i in eq. (3.7) through the procedure discussed in Section 3.3.1.2. The obtained values of the (time-invariant) weights k_i associated to each appliance are:

- 1) clothes dryer: $k_1 = 273$,
 - 2) kitchen fridge: $k_2 = 11$,
 - 3) dishwasher: $k_3 = 165$,
 - 4) heat pump: $k_4 = 444$.
- a calibration dataset, which consists of the measurements from May 30, 2012 to May 31, 2012. The calibration dataset is used to tune the hyper-parameters λ_1 and λ_2 in (3.7). Also in the calibration phase, the sub-metered power consumptions y_t^i are supposed to be available. The values of λ_1 and λ_2 are chosen through a cross-validation procedure, that is by minimizing (with a grid search) the *Total Relative Square Error (TRSE)* w.r.t. the calibration dataset, where the TRSE is defined as:

$$\text{TRSE} = \sum_{i=1}^N \frac{\sum_{t=1}^{T_c} (y_t^i - \hat{y}_t^i)^2}{\sum_{t=1}^{T_c} y_t^{i2}},$$

with T_c being the length of the calibration dataset. The chosen values of λ_1 and λ_2 are 10 and 750, respectively.

- a validation dataset D_T^v , which consists of the data for the days June 1-30, 2012 (as plotted in Figure 3.3). The proposed algorithm is applied to disaggregate the data of the set D_T^v .

To reduce the computational burden, a sub-optimal solution of Problem (3.7) is computed according to Section 3.3.1.2, by splitting the set of data to be disaggregated into 4320 subsets, each of equal length (i.e., 10 minutes). Finally, the aggregate power consumption observations are used to further reduce the computational complexity of Problem (3.7), by a-priori setting some parameters $x_t^{i,j}$ to 0. Specifically:

- at the time instants when the aggregate power consumption is less than 3000 W, the parameters $x_t^{i,j}$ associated to the clothes dryer and that multiply the basis $b^{i,j} = 4700$ W are set to 0;
- at the time instants when the aggregate power consumption is less than 1000 W, the parameters $x_t^{i,j}$ associated to the heat pump and that multiply the basis $b^{i,j} = 1900$ W are set to 0;
- at the time instants when the aggregate power consumption is less than 400 W, the parameters $x_t^{i,j}$ associated to the dish washer and that multiply the basis $b^{i,j} = 800$ W are set to 0.

Performance metrics The following metrics are used to assess the performance of the optimization-based algorithm:

3. Water and energy end-use disaggregation




Figure 3.3: Validation dataset D_T^V : aggregate electric power consumption from June 1, 2012 to June 30, 2012

- The *Estimated Energy Fraction Index (EEFI)*, defined as:

$$\hat{h}_i = \frac{\sum_{t=1}^T \hat{y}_t^i}{\sum_{i=1}^N \sum_{t=1}^T \hat{y}_t^i}.$$

The index \hat{h}^i provides the fraction of energy assigned to the i -th appliance, and it should be compared to the *Actual Energy Fraction Index (AEFI)*, defined as

$$h_i = \frac{\sum_{t=1}^T y_t^i}{\sum_{i=1}^N \sum_{t=1}^T y_t^i},$$

which in turn provides the actual fraction of energy consumed by the i -th appliance. The EEFI \hat{h}^i gives the users the information on how much energy each appliance is consuming, and so personalized hints for reducing their energy consumption can be provided.

- The *Relative Square Error (RSE)*, defined as:

$$RSE_i = \frac{\sum_{t=1}^T (y_t^i - \hat{y}_t^i)^2}{\sum_{t=1}^T y_t^i{}^2}.$$

The RSE provides a normalized measure of the difference between the actual and the estimated power consumption of the i -th appliance.

- The R^2 coefficient, defined for the i -th appliance as:

$$R_i^2 = 1 - \frac{\sum_{t=1}^T (y_t^i - \hat{y}_t^i)^2}{\sum_{t=1}^T (y_t^i - \bar{y}^i)^2},$$

with $\bar{y}^i = \frac{1}{T} \sum_{t=1}^T y_t^i$. Both the R^2 coefficient and the RSE measure how well the estimated power profiles match the actual power profiles over time. An accurate estimate of the power consumption profiles over time is essential to inform the customer on potential savings in deferring the use of some appliances to peak-off hours. Obviously, high value of the R^2 coefficients (or equivalently low values of the RSE) imply an accurate estimate of the EEFI.

Benchmark comparison: Factorial Hidden Markov Models The performance of the SOD algorithm presented in this paper is compared to the performance of the disaggregation approach based on *FHMMs* and implemented in the open source *Non-Intrusive Load Monitoring Toolkit (NILMTK)* Batra et al. (2014). For a fair comparison with the optimization-based algorithm presented here, the FHMM algorithm is trained based on the data from May 17, 2012 to May 31, 2012 and used to disaggregate the data belonging to the validation dataset D_v^T . The number of states of each HMM, or equivalently, the number of operating modes for each appliance, is the same across all appliances and it is set equal to 2.

3.4.1.2 Numerical results

The performance metrics introduced in the previous section and the estimated disaggregate power profiles are computed in order to assess the performance of the SOD algorithm, and benchmarking it against a 2-state FHMM algorithm. The obtained results are reported in Table 3.1, Table 3.2 and in Figure 3.4. It is worth remarking that the RSE and the R^2 coefficients, as well as the indexes \hat{h}_i and h_i , are referred to the portion of the dataset to be disaggregated (i.e., the whole month of June), while, for the sake of visualization, only a portion of the disaggregated power profiles is plotted in Figure 3.4.

The obtained results show that the developed optimization-based algorithm is able to accurately estimate the fraction of energy consumed by each appliance in the household (see Table 3.1). As a matter of fact, the EEFI \hat{h}_i is very close to the AEFI h_i for each appliance. This good performance is mainly due to an accurate estimate of the disaggregated consumption trajectories over the time (as shown by Table 3.2 and Figs. 3.4). The obtained results also reveal the following:

- both the optimization-based and the FHMM-based algorithms provide an accurate estimate of the power consumption of the clothes dryer. This is mainly due to the fact that clothes dryer events can be better distinguished from the other end-use events, as they usually show the highest power consumption peak and large durations;
- the FHMM-based algorithm slightly outperforms the optimization-based

3. Water and energy end-use disaggregation

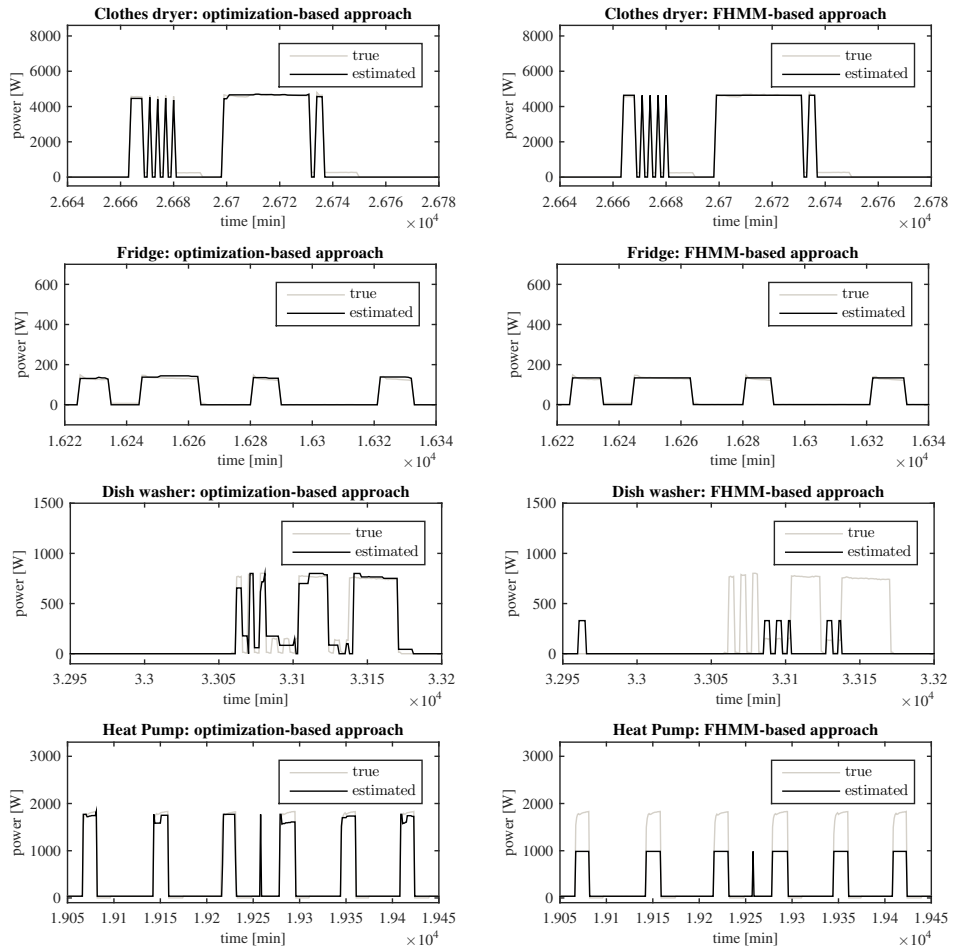


Figure 3.4: Disaggregate power consumption profiles. Results obtained through the optimization-based approach presented in the paper (left panels) and through the FHMM-based approach (right panels).

3.4. Applications to power load end-use disaggregation

Table 3.1: Fraction of energy assigned to each appliance (\hat{h}_i) and actual fraction of energy consumed by each appliance (h_i). Results obtained by using the optimization-based algorithm presented in the paper and the FHMM-based approach.

	SOD algorithm	FHMM-based algorithm	ground truth
Clothes dryer	30.7 %	31.6 %	31.3 %
Kitchen fridge	22.0 %	22.7 %	21.3 %
Dishwasher	4.0 %	7.7 %	5.1 %
Heat Pump	43.3 %	37.9 %	42.3 %

Table 3.2: Relative Square Errors and R^2 coefficients. Results obtained by using the optimization-based algorithm presented in the paper and the FHMM-based approach.

	SOD algorithm		FHMM-based algorithm	
	RSE _i	R _i ²	RSE _i	R _i ²
Clothes dryer	0.8 %	99.2 %	0.3 %	99.7 %
Kitchen fridge	24.2 %	63.3 %	20.6 %	68.7 %
Dishwasher	28.2 %	71.4 %	161.6 %	-63.9 %
Heat Pump	2.7 %	97.1 %	31.9 %	65.1 %

approach in the estimate of the fridge power consumption. This is mainly due to the fact that the power consumption profile of the fridge has a marked pattern, with periodic ON/OFF cycles, which is accurately captured by probabilistic models like Markov models;

- the optimization-based approach provides better performance than the FHMM-based algorithm in the estimate of the power consumptions of dishwasher and heat pump. In fact, the FHMM-based method tends to underestimate the consumption of the heat pump (see Figure 3.4, bottom panels) and thus to erroneously assign the residual power to the dishwasher.

We can conclude that the SOD algorithm is able to handle situations where multiple appliances are operating simultaneously, and also to accurately estimate the appliance power consumption profiles over time.

3.4.2 Power load disaggregation through HSID

3.4.2.1 Experimental settings

Data Accordingly to the experimental settings adopted for testing the sparse optimization-based algorithm, we tested the HSID algorithm against power load data from the AMPDs dataset. In order not to set a *a priori* number of appliances to consider for disaggregation, and at the same time avoid signal noise due to appliances with very low contribution to the total power consumption, we considered only the appliances contributing more than 5% of the total indoor consumption, i.e., heat pump, forced air furnace, clothes dryer, kitchen fridge, and security/network equipment. No outdoor uses, i.e., outside plugs, office uses and “hybrid/undefined” appliances, e.g., room aggregate consumption, were taken into account, in order to focus only on the most important residential activities and on traces measured at the end-use level for calibration, rather than those measured at the room level. We tested the supervised version of HSID against two sub-periods, extracted from the available 1-year dataset in order to account for possible seasonality effects on energy use:

- 6 Spring/Summer weeks data from May 16th 2012 to June 30th 2012, with the first two weeks used for appliances signatures extraction and FHMM calibration, and the remaining month for validation;
- 6 weeks from the Winter period, from November 16th 2012 to December 31st 2012, again divided into one third of the dataset for calibration and two thirds for validation.

These proportions are in line with those adopted in other state-of-the-art end-uses or energy conservation studies (e.g., Kolter et al., 2010; Kolter and Johnson, 2011; Farinaccio and Zmeureanu, 1999; Fischer, 2008). Since in this section we first apply the algorithm on a supervised case, we assumed data measured at the end-use level to be available for FHMM calibration purposes and signature extraction only during the 2-week training period, while for the validation period we used them only as ground-truth data for assessing the model outputs accuracy. Tests with the semi-supervised version of HSID do not consider the 2-week training period, but only single-event signatures for each appliance were assumed as input to FHMM and ISDTW. Both versions of the HSID algorithm parameters were set as follows: (i) L_e (event length for SDTW iteration) equal to 20 minutes; (ii) k (previous time step window for false positive event detection) equal to 30 minutes. Later in this chapter, we also assess the sensitivity of the algorithm with respect to different levels of signal noise and with respect to the number of appliances. For this latter experiment, we tested the HSID algorithm against the REDD dataset (Kolter and Johnson, 2011), considering a house which includes power consumption readings for up to 11 appliance types in the same household. Given the presence of missing readings and for consistency

with previous experiments, we down-sampled the original metering resolution (i.e., 3-5 seconds) to 1 minute, using 50% of the data (i.e., approximately 9 days) for calibration and the remaining 50% for validation.

Performance metrics The evaluation of the outcomes from disaggregation algorithms against a set of comprehensive and consistent metrics has been mentioned as one of the main challenges in the literature on NILM (Butner et al., 2013; Barker et al., 2014; Batra et al., 2014). We contribute an assessment of the quality of the disaggregation results from HSID against ground-truth data according to the following set of metrics, selected among the others because they overall cover the characteristics of the estimated end-use signals that should be considered for a complete evaluation.

- The *F-score* (F_s), as introduced in Batra et al. (2014), is evaluated for each appliance i according to the following formula

$$F_{s_i} = \frac{2 \times PC_i \times RC_i}{PC_i + RC_i} \quad (3.10)$$

where RC_i and PC_i are the *recall* and *precision*, respectively, evaluated for appliance i as $RC_i = \frac{TP_i}{TP_i + FN_i}$ and $PC_i = \frac{TP_i}{TP_i + FP_i}$. TP_i , FP_i , and FN_i are the number of events correctly classified when appliances are *on* (true positive), the number of events classified as *on* being the appliance actually *off* (false positive), and the number of events classified as *off* being the appliance actually *on* (false negative). The precision can be interpreted as a measure of how many detected events are relevant and the recall as a measure of how many relevant events are properly detected. The *F-score* indicator evaluates how good the algorithm is in classifying the operating states of the considered appliances. It ranges from 0 (0% accuracy on state detection) to 1 (100% accuracy on state detection).

- The assigned *PCE* gives information on the model accuracy in assigning the power consumption share to each appliance i , according to the following formula:

$$PCE_i = \frac{\left| \sum_{t=1}^H y_t^i - \sum_{t=1}^H \hat{y}_t^i \right|}{\sum_{t=1}^H \bar{Y}_t} \quad (3.11)$$

where y^i and \hat{y}^i are the ground-truth and estimated power consumption for appliance i respectively. An accurate algorithm would produce PCE values close to 0.

- The R^2 score, already defined in Section 3.4.1.1, assesses the accuracy of end-uses trajectories reproduction.

3. Water and energy end-use disaggregation

This set of metrics assesses the performance of NILM algorithms by different viewpoints corresponding to increasing levels of information provided on the signal characteristics and, correspondingly, different value for electric utilities and decision makers interested in designing energy demand management strategies.

F-score gives a basic level of information about the capability of NILM algorithms to properly detect the appliances operating states but does not provide any information about the power consumption. In turn, PCE gives an overall indication of the goodness of NILM models to estimate the power consumption assigned to each appliance. This is more informative than F-score to design customized feedbacks and other demand management strategies, since a model with a low PCE would be able to inform decision makers about major power uses and power ratios. Finally, the R2 score evaluates the finest aspects of NILM outputs, i.e., the accuracy in characterizing power consumption trajectories, which becomes essential for improving the information level on energy use. This allows to evaluate values of power consumption during peak periods, retrieve information about use frequencies and timing for major uses and, through the analysis of consumption pattern, identify changes in the electric equipment of a house or potential savings from equipment renewal.

Number of states for the FHMM component in HSID The computational complexity of FHMM algorithms grows exponentially in the number of states considered. It is in the order of $O(TM^{2N})$ for a problem with M states for each system element (appliance), N appliances and T time instances (Ghahramani and Jordan, 1997). On the other hand, the higher the number of states, the higher the number of appliances potentially extracted by the disaggregation, in principle. Yet, it is not easy to a-priori decide which number of states is suitable for accurately describing the consumption pattern of different fixtures just considering traditional FHMM, as each fixture has its own consumption pattern and *multi-state* or *continuously variable devices* might potentially require a very high number of states to be properly modeled. In order to explore how big such a number of states should be, we performed a preliminary sensitivity analysis by disaggregating a 1-month power consumption contributed by 4 appliances with distinct signatures that should maximize the potential of FHMM with many states.⁴ We considered an increasing number of states (2,3,4 and 7) and evaluated the variation in performance through the F-score and R2 metrics, which range from 0 to 1 in the best case (Batra et al., 2014). The values obtained for the two performance metrics are represented in the two radar plots in Figure 3.5, where each axis reports the performance metrics specifically for each

⁴Data from this experiment are retrieved from the AMPDs dataset (Makonin et al., 2013), similarly to other experiments in this chapter.

appliance and each coloured line connects the performance given by a specific setting of FHMM states. In principle, as the two metrics should be maximized, one would like to obtain a coloured line connecting all the vertices of the radar plot, with value 1.

Results show that the performance does not monotonically increase with the number of states, suggesting that increasing the number of states raises the computational cost of the algorithm, but this does not guarantee an improvement of the disaggregation accuracy. Consequently, alternative solutions to the increase of FHMM states need to be formulated in order to get accurate disaggregation results at reduced cost. Our HSID algorithm addresses this challenge from the point of view of computational complexity. In the FHMM module, the number of possible states allowed for each appliance to perform disaggregation is limited to 2, as mentioned in Section 3.3.2.1, thus representing the *on/off* state of each appliance. The FHMM module is then coupled with the ISDTW module explained in Section 3.3.2.1 to correct the end-use trace patterns. This combination of 2-state FHMM and ISDTW yields a complexity in the order of $O(T \times 2^{2N} + N \times \lceil T/L_e \rceil)$, where L_e is the time length defined for each event (therefore lower than T). From the point of view of disaggregation outputs, we expect the iterative correction of ISDTW on 2-state FHMM outputs to significantly improve the disaggregation accuracy.

3.4.2.2 Numerical results

Supervised HSID disaggregation Similarly to the experiments conducted with the SOD algorithm, we have comparatively analyzed the HSID algorithm with respect to the 2-state FHMM benchmark algorithm developed in Batra et al. (2014) (Figure 3.6). The radar plot in the figure must be read similarly to Figure 3.5 (the performance index relative to each appliance is reported on each axis and different colours refer to different algorithms) keeping in mind that, conversely to the F-score and R2 metrics, the PCE shows good performances for values close to zero.

The HSID algorithm overall achieves very good performance on all the three metrics considered. The F-score shows that the algorithm is able to correctly detect the operating states of each appliance with a rate higher than 95% on four appliances out of five simultaneously operating (and always higher than 70%). Slightly lower, but still very good, results are achieved also by the 2-state FHMM benchmark. Yet, the HSID algorithm significantly outperforms the benchmark on the other two metrics, attaining PCE values lower than 2% for all the five appliances and a R2 close to 1 for three out of four appliances. These results demonstrate that HSID is able to detect the operating states of the appliances, while also providing information on the contribution by each appliance to the total power consumption and on the consumption patterns for

3. Water and energy end-use disaggregation

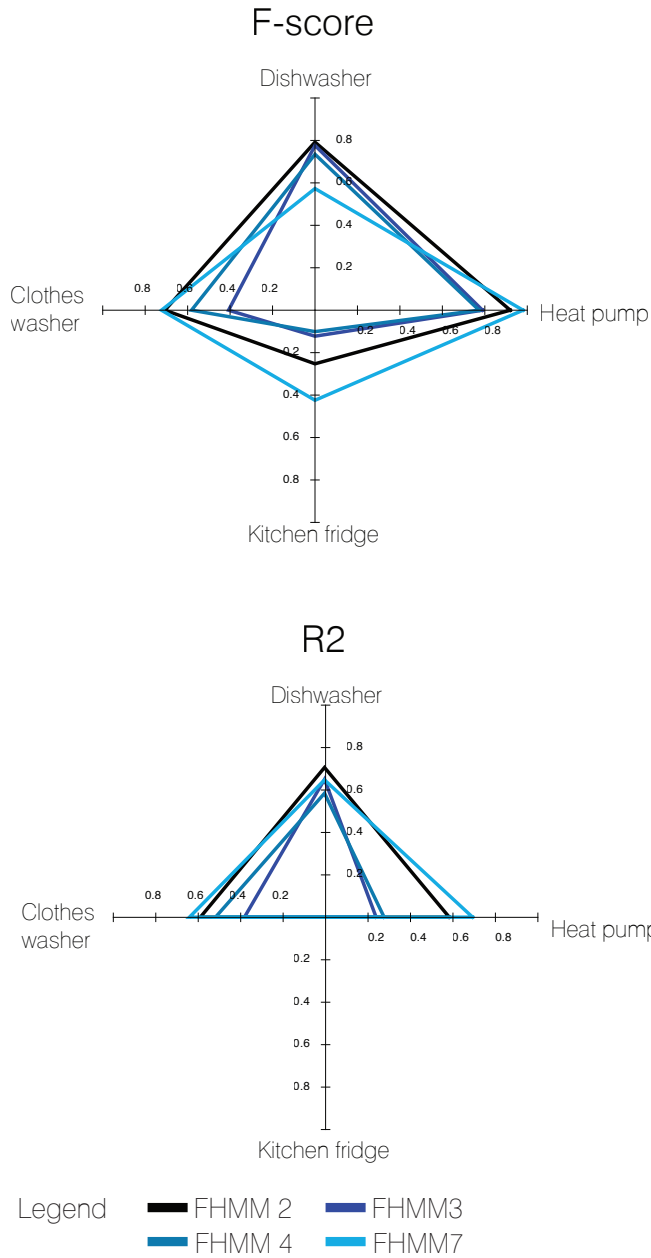


Figure 3.5: *F-score and R2 disaggregation accuracy on four appliances with increasing number of FHMM states.*

3.4. Applications to power load end-use disaggregation




Figure 3.6: Disaggregation performance metrics for supervised HSID and 2-state FHMM algorithm on 5 simultaneously operating appliances. From top: F-score, assigned PCE and R2 score. Data refer to the Summer period. The performance on the Winter period were found to be very similar.

3. Water and energy end-use disaggregation

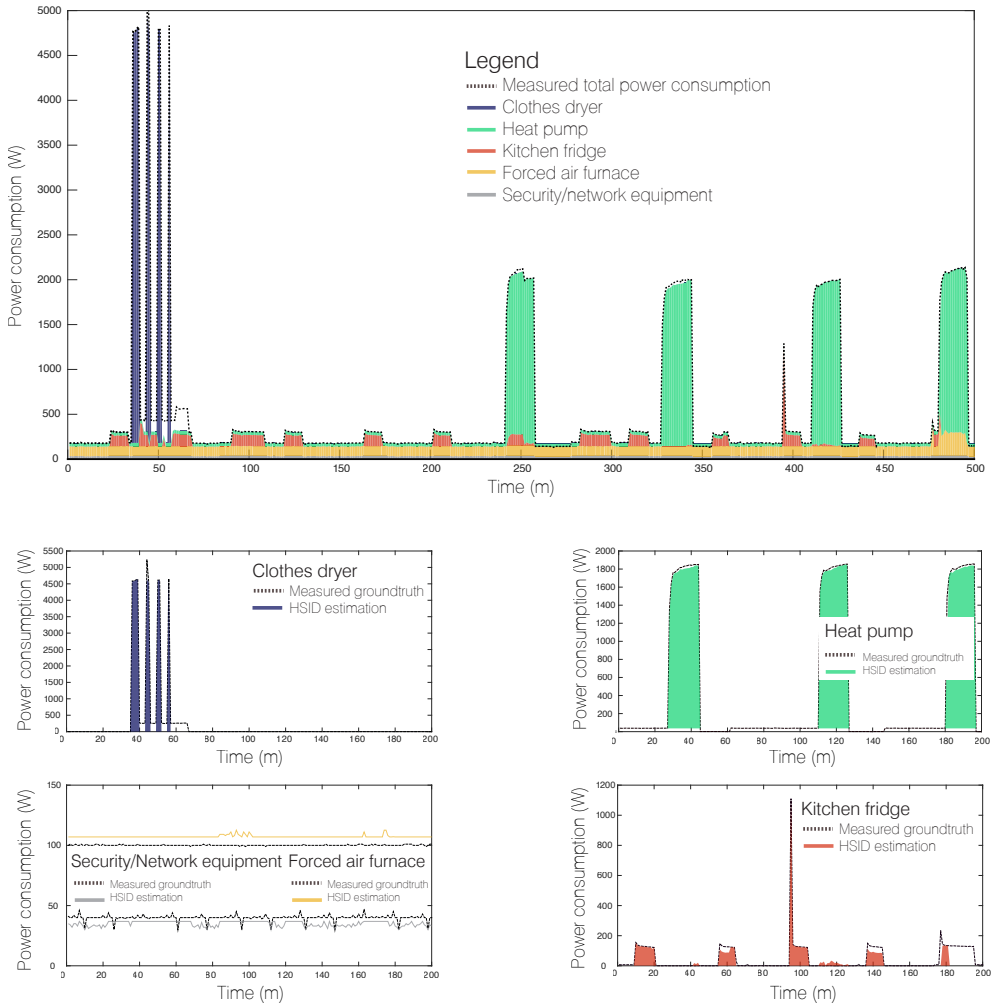


Figure 3.7: Consumption patterns estimated at the end use by the supervised HSID algorithm, compared with observed values. A stacked area with all the overlapping end-uses is represented in the top part of the figure. The four charts in the bottom part represent a detailed comparison of model output vs observed values of the 5 considered end-uses.

most of the appliances (Figure 3.7), despite multiple appliances operating simultaneously and, thus, overlapping end-uses. It is worth noticing that the largest improvement in correctly assigning power (Figure 3.6, middle panel) is mainly on the two most contributing appliances, thus the result is even more meaningful. Finally, another numerically relevant result is obtained on the estimation accuracy of the power consumption trajectories, which can be seen in the R2 metric (Figure 3.6, bottom panel) and by visual inspection of the trajectories contrasted against ground-truth data (Figure 3.7). While HSID performance, as measured by the three selected metrics, appears similar to SOD performance (even though slightly different performance metrics and input data from the same dataset were considered), Figure 3.7 shows that HSID outperforms SOD in accurately reproducing the shape of single-appliance power load trajectories. A careful analysis of Figure 3.7 also clarifies why for two appliances (i.e., the security equipment and the forced air furnace) the performance in terms of R2 is relatively small while the PCE is kept very low. The reason is that the trajectory of those two appliances is very noisy, thus being difficult to predict. However, it varies in a narrow range if compared to the other appliances', thus allowing for a correct estimation of its power contribution with average values. As mentioned in Section 3.3.2, the HSID algorithm does not perform signature correction for such appliances, as noise is prevailing and, therefore, trajectories cannot be reproduced accurately (even though visual analysis shows improvements with respect to the trajectories produced by the benchmark 2-state FHMM). Still, the algorithm is able to filter out the noise for such appliances and estimates average consumption for each of them as captured by the low values of PCE. The promising results we obtained for the supervised experiment showed the potential information content of appliances signatures, which appears to be key to allow for an accurate estimation of end-use power consumption trajectories. This finding, joint with the fact that HSID requires less training data than SOD and it is more computationally efficient than the latter, opens space for a discussion on the usability of the HSID algorithm under conditions where the availability of signatures is reduced, or the quality of the signal to disaggregate is low. For instance, we are interested in exploring HSID performance when just few end-use ground-truth data are available for algorithm training, or when the quality of the signal to disaggregate is altered, due for instance to measurement noise or the effect of many simultaneously operating appliances. We address both these questions with three *ad hoc* experiments in the next paragraphs. Firstly, the sensitivity of the algorithm to aggregate signal noise, as well as higher number of appliances, is assessed. Secondly, opportunities for exploiting the information content of signatures collected without intrusive measurements at the appliance level, while relying on the total consumption coupled with consumption diaries, are analysed and quantitatively discussed.

3. Water and energy end-use disaggregation

Algorithm sensitivity to signal noise The results discussed in the previous paragraph suggest that the HSID algorithm successfully attains good performance in supervised disaggregation problems, assuming the availability of a training dataset to accurately define the signatures of each appliance. The corrections introduced in HSID through signature matching allow to accurately approximate end-use consumption trajectories. Yet, no signal noise was considered, as the total electricity consumption trace to disaggregate consisted of the sum of appliance-level traces. In real-world applications, the quality of the signal as measured by smart meter can be affected by noise, due to measurement errors or interference among appliances. We expect this noise to affect the quality of the approximated disaggregated traces and, in order to evaluate the impacts of less accurate signatures on the disaggregation performance, we ran the following sensitivity analysis. First, we estimated the measure error from AMPDs dataset as the difference between the total power load trace measured by the house-level smart meter and the sum of appliance-level power load traces. Afterwards, we repeated the supervised disaggregation experiment commented in the previous section upon summation of that noise to the total power trace to disaggregate. More specifically, we considered a noise ratios of 0.5, 1, 1.25, and 1.5 for each algorithm run, in order to evaluate the sensitivity of the HSID performance with respect to different noise magnitudes. The results of this sensitivity analysis are represented in Figure 3.8.

The algorithm shows a good robustness under conditions of noisy signal over the F-score metric, as the introduction of signal noise affected significantly only the disaggregation accuracy of the kitchen fridge trace, but it is still kept to values around 0.6. As expected, signal noise caused a performance degradation for those metrics more strictly linked to trace shape and pattern. Indeed, PCE increases to values higher than 5% for all appliances, as soon as noise is considered and R2 drops to zero for the kitchen fridge appliance. Still, PCE is lower than 20% for all appliances when the actual noise of the dataset is introduced (i.e., noise ratio equal to 1) and it is lower than 40% for all appliances in the worst case when the actual noise is artificially increased by 50% (i.e., noise ratio 1.5), being overall lower than 6% for 3 appliances out of 5, in all noise scenarios considered. Also the performance in terms of R2 is lowered as the noise level increases, as signal noise makes it harder to reproduce accurate trajectories through signature matching. Still, for noise ratios lower than 1.25 its values are still higher than 0.6 for heat pump and clothes dryer, i.e., two out of the three appliances that presented positive values of R2 in the supervised experiment without signal noise. We can conclude that the disaggregation performance of HSID can be significantly affected by signal noise, but overall the algorithm shows good performance robustness in terms of event detection (F-score) and appliance-level contributions for noise levels in the range of the ac-

3.4. Applications to power load end-use disaggregation




Figure 3.8: Disaggregation performance metrics for HSID under different levels of noise on the smart-metered trace to disaggregate. From top: F-score, PCE and R2 score.

3. Water and energy end-use disaggregation

tual measured noise.

Algorithm sensitivity to higher number of appliances The results discussed in the previous paragraphs demonstrate that the HSID algorithm is able to accurately disaggregate the end-use consumption of the major appliances (i.e., devices contributing more than 5% of the total power consumption), also showing sufficient robustness against increasing levels of noise. In this section, we assess the sensitivity of the algorithm’s performance when tested against a larger set of appliances. Specifically, we test the disaggregation on one household from the REDD dataset (Kolter and Johnson, 2011), which includes 11 different appliance types. The results of this experiment, computed on the validation dataset, are reported in Table 3.3, with the algorithm performance measured for each appliance in terms of the three metrics formalized in Section 3.4.2.1.

Despite HSID is trained on a shorter period than the previous experiment, it attains good performance in terms of F-score metric, with an average value across the 11 appliances equal to 0.61. More precisely, the F-score metric is greater than 0.8 for 6 out of 9 appliances (and equal to 1 for 4 appliances), and only 4 out of 11 appliances are below 0.5. This high disaggregation accuracy is confirmed by the values of PCE. HSID attains an average value in this metric below 4%, with the maximum error equal to 14% for the disaggregation of the lighting power consumption, thus outperforming the FHMM disaggregation method of Kolter and Johnson (2011) which reported an average error of over 50%. Finally, the performance evaluated in terms of R2 shows the largest degradation, with positive values attained for only 3 appliances. This low performance can be explained by the increased variability in the total metered consumption when numerous heterogeneous appliances are considered and overlap. It is worth noting that HSID attains positive values of R2 for both the dishwasher and the fridge (included in the kitchen outlets), which are the appliances contributing the most to the total consumption as well as the ones characterized by a regular signature. These results suggest that increasing the number of appliances impact negatively on the ability of HSID of reproducing the single end-use trajectories, as shown by the low values of R2. However, the algorithm attains good performance in the other two metrics, thus showing a good scalability to large set of appliances for both event detection and recognition of appliance-level contributions.

Semi-supervised HSID disaggregation In order to test the performance of the semi-supervised extension of HSID as described at the end of Section 3.3.2.2, we run an experiment similar to the one performed for the supervised version of HSID, but considering that only a single-event signature for each

3.4. Applications to power load end-use disaggregation

Table 3.3: Disaggregation performance metrics for HSID tested on a house comprising 11 appliances from the REDD dataset. Fields filled with "-" represent R2 values lower than 0.

	F-score	PCE (%)	R2
Kitchen outlets	1	8.12	0.17
Lighting	1	14.15	-
Furnace	1	2.80	-
Dishwasher	0.61	0.73	0.81
Stove	1	2.74	0.03
Washer-dryer	0.05	8.96	-
Miscellaneous	0.99	0.01	-
Bathroom GFI	0.82	1.96	-
Outlets unknown	0.20	1.96	-
Air conditioning	0.03	0.81	-
Smoke Alarms	0.01	0.13	-

appliance is available, as retrieved from the total power consumption pattern. Figure 3.9 reports the results of this semi-supervised application.

Not surprisingly, results show that the semi-supervised use of the HSID algorithm is underperforming with respect to the supervised, partially intrusive experiment. However, some important and promising aspects do emerge. First, despite the small amount of information available, the performance is still acceptable: only one appliance has an F-score lower than 0.7, all appliances have PCE lower than 10%, and still two appliances present trajectories estimation accuracy (R2) higher than 0.8. Second, it is worth examining the reason of this degradation of performance, which is probably twofold. The one-event signatures produce a loss of information on the potential variability of the appliances signatures, which are instead represented in the 2-week training dataset used in the supervised application. Moreover, the one-event signatures fail by definition in characterizing the behavioural component of the power consumption, i.e., the information about energy use habits that can only be *a-priori* retrieved through the study of end-use consumption patterns for a significantly longer period. We took into account such information in the supervised application by the FHMM module, where the data retrieved from an intrusive measurement period allowed for training the probabilities of initial operating state and state transition for each appliance. These probabilities constitute the statistical expression of energy consumers' habits. In contrast, this information is missing in the semi-supervised experiment, causing a performance decline. However, the results of this semi-supervised application of the HSID algorithm suggest that the contribution given by the use of signature patterns is essential to accurately

3. Water and energy end-use disaggregation




Figure 3.9: Disaggregation performance metrics for semi-supervised HSID. From top: F-score, assigned PCE and R2 score.

perform power load disaggregation.

3.4.3 Multi-resolution power load disaggregation through supervised HSID

In addition to the tests described in this section, we performed a series of end-use disaggregation experiments with HSID algorithm against data with progressively down-sampled resolutions, in order to assess the effect of smart meter data sampling resolution on end-use disaggregation capabilities and information loss with respect to the best available resolution and, ultimately, understand possibilities for an actual usability of the disaggregation algorithms we developed on data with a sampling resolution of 1 hour. As highlighted among smart metering data gathering challenges (see Section 2.2 in Chapter 2), the motivation behind such analyses is that several utilities in Europe and worldwide are collecting data with lower resolutions than a few seconds or minutes because of customers' privacy, as well as technological and data storage issues (e.g., smart meter battery life duration, data transfer and storage costs).

3.4.3.1 Experimental settings

Data Coherently with the experiments described in Section 3.4.2, we considered the AMPDs dataset and, again, the power load contributions given by those appliances contributing more than 5% of the total power load consumption, i.e., heat pump, forced air furnace, clothes dryer, fridge and security/network equipment. We considered a consumption period of 45 days during the months November/December 2012, with 2-week data for calibration and a month of validation. The training dataset consists of the trajectories of each end-use and their sum over the training period.

Data resolution downsampling We down-sampled the original 1-minute power load trajectories to build trajectories with the following sampling resolutions: 5, 15, 30, 60, 120 minutes.

Performance metrics We adopt the same performance metrics defined in Section 3.4.2.1 i.e., F-score, PCE, and R2 score are considered. The groundtruth benchmark for performance assessment is the validation dataset sampled for each single appliance at 1-minute resolution, i.e., the finer resolution available.

3.4.3.2 Numerical results

Performance metrics obtained for HSID on multi-resolution power load end-use disaggregation are reported in Table 3.4, Table 3.5, and Table 3.6.

3. Water and energy end-use disaggregation

Table 3.4: *F-score for supervised HSID applied to power load data with different data sampling resolutions.*

	Data sampling resolution					
	1 min	5 min	15 min	30 min	60 min	120 min
Heat pump	0.95	0.80	0.38	0.30	0.25	0.23
Forced air furnace	0.99	0.99	0.99	0.99	0.99	0.99
Clothes dryer	0.80	0.25	0.28	0.11	0.15	0.09
Kitchen fridge	0.95	0.57	0.27	0.49	0.50	0.50
Security/	1	1	1	1	1	1
Network equipment						

Table 3.5: *PCE (%) for supervised HSID applied to power load data with different data sampling resolutions.*

	Data sampling resolution					
	1 min	5 min	15 min	30 min	60 min	120 min
Heat pump	1.3	2.2	20.3	20.8	21.4	5.3
Forced air furnace	1.6	2.3	2.1	1.1	0.7	0.007
Clothes dryer	0.06	10	9.6	5.7	4.0	0.6
Kitchen fridge	0.7	2.2	2.8	0.4	0.5	1.4
Security/	1.2	0.2	0.03	0.02	4E-5	0.02
Network equipment						

Table 3.6: *R2 for supervised HSID applied to power load data with different data sampling resolutions. Fields filled with "-" represent R2 values lower than 0.*

	Data sampling resolution					
	1 min	5 min	15 min	30 min	60 min	120 min
Heat pump	0.98	0.66	0.04	-	-	0
Forced air furnace	0.41	-	-	-	-	-
Clothes dryer	0.99	0.08	-	-	-	-
Kitchen fridge	0.77	-	-	-	-	-
Security/	-	-	-	-	-	-
Network equipment						

3.4. Applications to power load end-use disaggregation

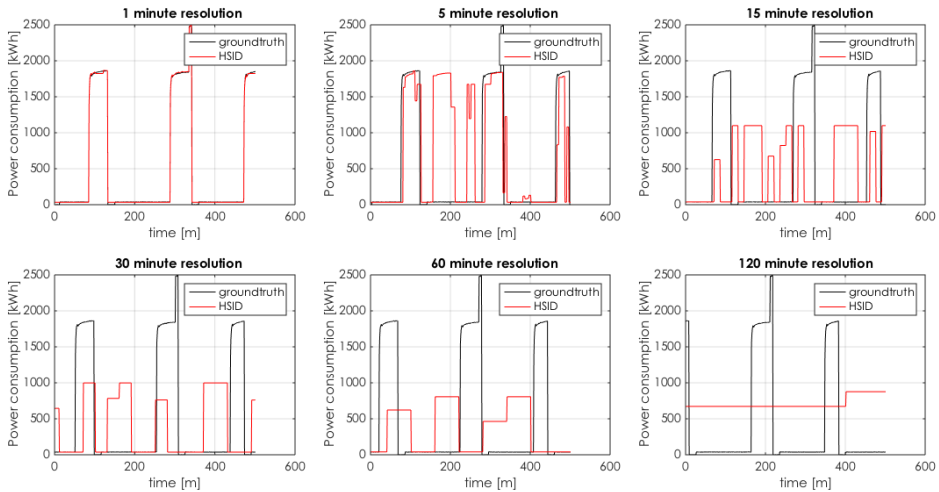


Figure 3.10: Heat pump groundtruth trajectory (black line) and trajectory estimated by HSID (red line) at different resolutions.

It clearly appears that, for most of the considered appliances, F-score and R2 quickly decrease for data sampling resolution lower or equal to 5 minutes. This happens also for those appliances that have the largest overall contribution, the biggest operating range and a marked signature, like the heat pump, which is the most contributing appliance and whose groundtruth and estimated trajectories at different resolutions are shown in Figure 3.10. The only exception holds for the forced air furnace and the security/network equipment, which keep a high F-score. However, this is attributable to the fact that they are operating for most of the time with values close to their average, but they are also very noisy, as demonstrated by the low R2 and its sudden drop. Different comments can be drawn for the power contribution error metric, which gets worse for resolutions of 5 to 15 minutes, but then in some cases (e.g., kitchen fridge, clothes dryer) tends to improve again as the resolution lowers down to 120 minutes. Although unexpected, the most likely reason for that improvement is that, with lower resolutions, the estimated signal tends to become closer to a constant signal representing the average end-use consumption over time for the specific appliance (see bottom-right plot in Figure 3.10).

As a consequence, if the training data well represent the average use of a specific appliance, the integral under the estimated power load trajectory over time is likely to be close to the actual consumption of that appliance. Yet, as the power contribution error metric measures the estimate error on the aggregate power consumption, it hides the low accuracy in terms of on/off event detection and accuracy in reproducing the actual consumption trajectories. This represents a strong limitation to disaggregation capabilities at low resolutions. Indeed,

3. Water and energy end-use disaggregation

even though apparently lower resolution than the minute still allow for an accurate estimate of the overall contribution of each end-use, firstly this accuracy depends on the meaningfulness of the training period, that requires intrusive measurements. Secondly, the outcome of the disaggregation process should be accurate in terms of trajectory estimate accuracy (i.e., F- score and R2 score) in order to be fully informative to the demand side management. An accurate estimation of the end-use consumption trajectories is essential to provide information on timing of appliance use during the day, peak demands and frequencies of use.

The results obtained from the above described experiments suggest that resolutions higher than few minutes strongly affect the ability to retrieve such information. Despite the limited testing does not allow a complete generalization of this evidence, the result obtained is coherent with the work by Carrie Armel et al. (2013), who claim that most of the end uses can accurately be identified with resolutions of 1s-1m, while the number of identifiable appliances strongly reduces and only a very small number of main appliances can be identified at 1-hour resolution by visually observable patterns or correlation with external variables (e.g., temperature), which does not mean they can be accurately disaggregated with automated end-use algorithms.

3.5 Water end-use disaggregation through HSID

In the previous section, we showed promising results obtained with the application of our two disaggregation algorithms (i.e., SOD and HSID) to power load disaggregation. In particular, we demonstrated that the second algorithm, i.e., HSID, performs well on the disaggregation of 1-minute resolution data, even with noisy signals and increasing number of appliances, and can accurately estimate single-appliance contributions even for slightly lower resolutions (e.g., 5 minutes). In this section, we explore the portability of HSID to water end use disaggregation.

In the water sector, the implementation, calibration, and validation of most end-use water demand models have been, so far, relying on end-use data collected with ad-hoc experimental trials and research projects, involving a limited number of households, usually lower than 100 (e.g., Mayer et al., 2004; Froehlich et al., 2009; Suero et al., 2012). The data collected within these initiatives are hardly shared or made public available because of privacy and data access issues. Given the unavailability of complete public datasets with high-resolution water end-use data for algorithm calibration and testing, in this section we first describe an open-source, stochastic simulation model that we developed for synthetically generating high-resolution time series of water use at the end-use level. We then discuss the portability of HSID to water end-use disaggregation

through practical applications onto the data stochastically generated.

3.5.1 Developing a stochastic simulation model for the generation of residential water end-use demand time series

In the model we propose here, each water fixture is characterized by its *signature* (i.e., its typical single-usage pattern (Cardell-Oliver, 2013b)), along with the frequency distributions of its number of uses per day, single use duration, time of use during the day, and contribution to the total water demand. Such statistics, aggregated from high-efficiency water demand data gathered in 2005-2006 from over 300 single family homes located in 9 cities across the USA, are conditioned to the number of occupants of the household, the presence of water consuming fixtures, and appliance efficiency levels, which represent the only inputs required to run the model for a single household. At the best of authors knowledge, only few works provide tools for synthetically generating end-use water demand data through stochastic generation and Monte Carlo sampling (e.g., Blokker et al., 2010; Rosenberg et al., 2007; Abdallah and Rosenberg, 2014; Escriva-Bou et al., 2015a, , or the online tool Waterville⁵). Our model advances these tools as (i) it is built upon a dataset of high-resolution, smart-metered water demands, consistently related to a set of appliance statistics from the same household sample, (ii) it generates end-use time series of water use preserving the typical usage pattern of each fixture, thus end-use events do not consists of simple demand pulses, and (iii) it allows simulating household water demands under diverse social and technological scenarios, as well as sampling resolutions and sample size, thus allowing the generation of scenarios different form the one represented by a direct analysis of end-use calibration data.⁶

3.5.1.1 Method

Our stochastic simulation algorithm operates according to the phases detailed below. In short, given a sample of houses with associated number of occupants, available water consuming fixtures, and fixture efficiency, the model simulates both the end-use time series of water use and their aggregation as total household water demand, with a time resolution (i.e., the simulation time step) up to 1 second, for the defined simulation horizon.

Assumptions Considering a single house, in our stochastic model we assume that the time series of water use of a generic j -th water consuming fixture (e.g., toilet, faucet, shower, etc...) during a given d -th day can be characterized by the following elements:

⁵<http://waterville.hrwallingford.com/waterville/>

⁶Scenario generation is not analyzed in this thesis, but a few exaples can be found in Cominola et al. (2016a)

3. Water and energy end-use disaggregation

- number of times the j -th fixture is used during the day (each time it is used will be later referred as *consumption event*);
- starting time of use during the day for each consumption event;
- duration of each consumption event, in seconds;
- volume of water required by each consumption event, in liters.

In addition, we assume that each end-use consumption event can be shaped according to the specific *signature* of its fixture, i.e., the characteristic water use series over time for a single consumption event of a specific end-use, such as the ones represented in Figure 3.11, properly stretched according to consumption event duration and water volume. A number of stretched signatures equal to the number of daily consumption events of each fixture are then assembled to form the end-use daily time series of water use. Finally, we assume that each fixture can be operated at most only once during each simulation time step.




Figure 3.11: Sample signatures for standard-efficiency toilet and shower fixtures.

Model inputs Our stochastic model requires the specification of the following inputs:

- sample size N , i.e., number of households for which the model will simulate water demands;
- house size $H = \{h_1, h_2, \dots, h_N\}$, $h_i > 0 \forall i \in [1, N]$, i.e., number of occupants for each house in the sample;
- fixture presence $K_i = \{k_1, k_2, \dots, k_M\}$, $k_j \in \{0, 1\} \forall j \in [1, M]$, i.e., a binary index specifying the presence (absence) of the j -th fixture in the i -th household;

- fixture efficiency level $E_i = \{e_1, e_2, \dots, e_M\}$, $e_j \in \{0, 1\} \forall j \in [1, M]$, i.e., a binary index specifying the efficiency level (*standard* or *high*) of each fixture in each household. Standard and high efficient fixtures are assumed to differ as for the water demand and duration required by their usage, as well as their signature. All other features (e.g., time of use) are assumed to be commonly shared by normal and efficient fixtures.
- length of the simulation horizon T , in days.
- time sampling resolution r , $r > 0$ for the output demand trajectories, in seconds.

End-use water demand traces generation Given the input listed in the previous paragraph, let's consider the i -th house, characterized by h_i occupants, fixture presence K_i and fixture efficiencies E_i . Our stochastic model relies on a database containing M signatures (one for each fixture) and a database containing frequency distributions of the number of uses per day, event durations, water volumes, and time of use during the day for each fixture. Each distribution is conditioned to the number of house occupants (h_i) and fixture efficiencies ($e_{i,j}$). From this information, the model generates time series of water use according to the following phases:

1. **Number of consumption events.** The model samples the number of consumption events for each fixture j and each day d of the simulation horizon T as $NCE_{i,d,j} \sim F(NCE_j | h_i, e_{i,j})$, where $F(NCE_j | h_i, e_{i,j})$ is the frequency distribution of the number of usages per day for appliance j , conditioned to the number of house occupants (h_i) and fixture efficiencies ($e_{i,j}$).
2. **Consumption event statistics.** For each water consumption event $l \in [1, NCE_{i,d,j}]$, the model then samples the following elements:
 - event duration $D_{i,d,j} \sim F(D_j | h_i, e_{i,j})$;
 - water volume $V_{i,d,j} \sim F(V_j | h_i, e_{i,j})$;
 - time of use $TOU_{i,d,j} \sim F(TOU_j | h_i, e_{i,j})$.

$F(D_j | h_i, e_{i,j})$, $F(V_j | h_i, e_{i,j})$, and $F(TOU_j | h_i, e_{i,j})$ are, respectively, the frequency distribution of consumption event duration, water volume, and daily time of use for the j -th fixtures, conditioned to h_i and $e_{i,j}$. In the current version of the model, the frequency distributions for event duration, water volume, and time of use are assumed to be mutually independent.

3. **Consumption event trace creation.** The time series of water use of each water consumption event is generated by stretching the specific signature of the considered fixture according to the sample values $V_{i,d,j}$ and $D_{i,d,j}$.

3. Water and energy end-use disaggregation

In order to do so, first randomly chosen points of the signature are iteratively removed/replicated, in order to match the desired event duration $D_{i,d,j}$. Then, the magnitude of each point of the signature is scaled so that the integral over the signature matches the desired water volume $V_{i,d,j}$. After a signature is stretched according to the above procedure, it is positioned over the end-use time series of water use $y_{i,j}$ according to $TOU_{i,d,j}$.

This procedure is iterated from step 1 to step 3 until the simulation is completed, for all the M fixtures and the simulation period T .

Model output Our stochastic end-use model returns, as output, the end-use time series of water use $y_{i,j}$ for each house i and its fixtures j , as well as the total household water demand $Y_i = \sum_{j=1}^M y_{i,j}$. We showcase a sample outputs of the stochastic end-use simulation model in Figure 3.12. The generated traces show the heterogeneity of the different fixtures' signatures, with the toilet and the faucet characterized by an almost instantaneous pulse, while shower and clothes washers correspond to longer consumption events (Figure 3.12, left panel). Moreover, it is worth noting the differences in peak flow, duration and pattern of the two faucet consumption events, which are strongly dependent on the multipurpose manual usage of this fixture. The hourly distribution of the consumption events during a sample day (Figure 3.12, right panel) shows a typical two-peaks pattern, with the shower and the dishwasher contributing to the morning and evening peak, respectively. The use of toilet and faucet is more equally distributed during the day, while the clothes washer is simulated overnight as this use was observed to occur during off-peak times with positive frequencies.

Data We retrieved the signatures of water consuming fixtures, as well as their associated statistics from DeOreo (2011) and Abdallah and Rosenberg (2014). In the aforementioned study, Aquacraft Inc. metered end-use data from 280 standard houses and 25 high-efficiency houses across 9 U.S. cities, for a period of over 2 weeks between 2005 and 2009 at 10 seconds resolution. Overall, Aquacraft metered 753,076 events during 4,036 days.

It is worth mentioning the following hypotheses we used in the implementation of our stochastic model:

- We only considered indoor water uses, including the following fixtures: toilet, clothes washer, shower, dishwasher, faucet. Toilet and clothes washer signatures and statistics distinguish between standard and high-efficiency, while no distinctions based on efficiency is present for the other appliances.

3.5. Water end-use disaggregation through HSID




Figure 3.12: Samples of end-use trajectories output of the stochastic model for a single household over few hours (left panel) and one day (right panel). For the sake of visualization, the data are aggregated at one minute resolution (left panel) and one hour resolution (right panel).

- We modeled houses with 1, 2, 3, 4, 5, > 5 occupants. Houses with more than 5 occupants are grouped together in the last category;
- Water fixtures signatures are directly extracted from Acquacraft’s report (DeOreo, 2011) using the Get Data Graph Digitizer software⁷ and rescaled to 1 second resolution, according to the procedure described in Gaiardelli (2015). Moreover, we computed a median signature for each of those appliances that allowed the extraction of multiple signatures.
- Water fixture statistics (i.e., frequency distributions of number of uses per day, consumption event durations, water volume and time of use during the day) are derived, for each appliance, number of house occupants, and level of efficiency, from the data gathered in DeOreo (2011).

3.5.2 Multi-resolution water disaggregation through supervised HSID

3.5.2.1 Experimental settings

Data We generated total and end-use water consumption trajectories for 45 days (15 for HSID calibration, 30 for validation), using the stochastic data generator described in Section 3.5.1, and considering end-use probability distributions for a house with 2 occupants (i.e., the house size with higher frequency in the input dataset). Coherently with the resolution of power load data used in the first part of this chapter, we generated that dataset for five different water consumption fixtures (toilet, shower, faucet, dishwasher, clothes washer) with 1-minute sampling resolution as a baseline and benchmark for validation.

⁷Software available at <http://getdata-graph-digitizer.com/>, last visited on 30/03/2016.

3. Water and energy end-use disaggregation

Table 3.7: *F-score for supervised HSID applied to water data with different data sampling resolutions.*

	Data sampling resolution		
	1 min	5 min	60 min
Toilet	0.46	0.53	0.14
Shower	0.57	0.24	0.07
Faucet	0.44	0.39	0.26
Dishwasher	0.33	0.16	0.05
Clothes washer	0.22	0.09	0.01

Table 3.8: *PCE (%) for supervised HSID applied to water data with different data sampling resolutions.*

	Data sampling resolution		
	1 min	5 min	60 min
Toilet	3.6	4.9	12.6
Shower	5.8	4.6	8.1
Faucet	2.2	8.7	13.1
Dishwasher	7.8	5.2	7
Clothes washer	9.6	0.3	7

Data resolution downsampling In order to test HSID performance both onto high-resolution data (i.e., 1 minute), as well as lower ones, we down-sampled the data to 5 and 60 minutes, as 5 minutes sampling resolution resulted to be critical from the multi-resolution disaggregation experiments (see Section 3.4.3) on power load data and 60 minutes represents the sampling resolution commonly adopted in several locations worldwide.

Performance metrics Coherently with the experiments performed on power load disaggregation, we measured HSID performances in terms of F-score, PCE and R2 score (see Section 3.4.2.1 for metrics definitions).

3.5.2.2 Numerical results

In general, the performance (see Tables 3.7 - 3.9) is much lower than the one obtained for disaggregating the same number of appliances considering power load data, already at 1-minute resolution. We attribute this drop in performance, especially for what concerns operating status detection (F-score) and accuracy in reproducing water consumption trajectories (R2), to the following

Table 3.9: *R2 for supervised HSID applied to water data with different data sampling resolutions. Fields filled with "-" represent R2 values lower than 0.*

	Data sampling resolution		
	1 min	5 min	60 min
Toilet	0.20	0.20	0
Shower	0.51	0.03	0
Faucet	-	0.14	0.07
Dishwasher	-	-	-
Clothes washer	-	-	-

reasons. As Figure 3.13 shows, if we plot the total water consumption trajectory at 1-minute resolution, the operating range of water consumption appliances is narrow and no clearly identifiable signatures are present. This clearly makes the discrimination among end-uses harder and lowers the effect of HSID signature corrections. In addition, as it can be seen again in Figure 3.13 and also in Nguyen et al. (2013a), many water appliances operate with sub-minute cycles or human-operated events, therefore downsampling really makes it hard to distinguish between different appliances (e.g., a toilet flush and a short activation of a faucet look similar at the minute resolution and their consumption pattern reduces to be one point).

Moreover, the following comments already reported for the electricity case study, are confirmed: (i) the accuracy in reproducing the trajectory of the most contributing appliances (i.e., toilet and shower) is slightly higher than that for the other appliances, and (ii) downsampling from 5-minutes to 60-minutes significantly worsen the disaggregation accuracy.

In accordance to the results shown in the previous section, also in the case of multi-resolution water consumption disaggregation the performance in terms of power contribution error remains acceptable even at 1-hour sampling resolution (the maximum error is approximatively 13% on the faucet). However, as mentioned in the previous section, a good performance measured by the power contribution error metric does not reflect into the accuracy in end-use trajectory reproduction, which is needed in order to understand frequencies and appliance time-of-use information, ultimately useful to inform the design of customised water demand management strategies. In conclusion, the results obtained so far suggest that an accurate end-use disaggregation cannot be performed with the available disaggregation algorithms when the data sampling resolution is lower than a few minutes, or even better, seconds.

3. Water and energy end-use disaggregation




Figure 3.13: *Sample trajectory of 1-minute total water consumption data, generated for a 2-occupants house.*

3.6 Discussion

In this chapter we presented two algorithms for water and energy end-use disaggregation. The first algorithm, SOD, is based on the assumption that the unknown appliance power/water consumption profiles are piecewise constant over time and it exploits the information on the time-of-day probability in which a specific appliance might be used. The disaggregation problem is formulated as a least-square error minimization problem, with an additional (convex) penalty term aiming at enforcing the disaggregate signals to be piecewise constant over the time. The second algorithm, called Signature-based Iterative Disaggregation HSID, is based on the combination of Factorial Hidden Markov Models, which provides an initial approximation of the end-use trajectories, and Iterative Subsequence Dynamic Time Warping, that processes the end-use trajectories in order to match the typical power/water consumption pattern of each appliance.

We assessed the performance of the two algorithms onto real-world power consumption data with 1-minute sampling resolution, and benchmarked them

against state-of-the-art algorithms, showing that they both achieve high disaggregation accuracy. Moreover, we demonstrated that HSID accurately reproduces the consumption trajectory of each end-use, which is a major challenge in the field of energy disaggregation since many NILM algorithms can only afford the detection of operating states. The accuracy demonstrated by HSID in reproducing the consumption trajectories overcomes this challenge and allows supporting the design of more informed demand management strategies on the basis of the additional information on single appliance contribution to total consumption, use frequencies, and timing, which are essential to estimate and control peak and base load demand. HSID was also demonstrated to be robust with respect to noisy signals, scalable to dataset including a large set of appliances and can be successfully used in non-intrusive experiments without requiring appliance-level measurements. Finally, we extend the application of HSID to the disaggregation of synthetic water data at multiple data resolutions.

Yet, algorithms performance progressively drops for power load data with lower resolution than 1 minute, and significantly lowers when considering water consumption data. The latter are, in fact, characterised by consumption events with shorted duration, often in the order of few seconds or few minutes, and with appliances operating in similar ranges, which makes it harder to extract their consumption pattern from the aggregate, total consumption, signal. This result is coherent with the literature on energy end-use disaggregation Carrie Armel et al. (2013) and with the fact the the few works developed in the literature on water end-use disaggregation consider data metered with sub-minute resolution Nguyen et al. (2013a). Still, a promising result from the multi-resolution disaggregation on water data is that the overall contribution of each end-use can be accurately estimated even at low resolution and, therefore, the ranking of most consuming fixtures can be retrieved with an acceptable error. However, this would require a suitable and representative calibration dataset, collected intrusively by directly metering each appliance, or with an extensive campaign of water consumption diaries collection. The outcomes found in terms of suitable sampling resolution and algorithms development can support electric and water utilities and planners in the choice of suitable meters (i.e., with suitable data sampling), including data sampling resolution as a criterion for the choice, together with data storage and transmission resources and costs.

In general, the obtained outcome represents an opportunity for further development of NILM methods in a truly non-intrusive perspective, especially if supported by social computing, consumers' involvement techniques and ICT platforms (Giuliani and Mossina, 2016). Indeed, platforms for increasing consumers' awareness and engagement in demand management are undergoing a period of significant development, representing promising tools to pursue the

3. Water and energy end-use disaggregation

usability and scalability algorithms like HSID at reduced cost.⁸ Ongoing and future research should focus, at a first stage, on a further testing of the algorithms against other datasets, possibly gathered in different spatial contexts and with different meter resolutions. Also, potential use for completely non-intrusive applications should be further explored and tested on real cases, coupled with the design, implementation and monitoring of demand-side management strategies (Gaiser and Stroeve, 2014).

⁸See, for instance the Opower 3.0 platform at http://opower.com/company/news-press/press_releases/10

4

Descriptive and predictive modelling of residential water users

Abstract¹

The effectiveness of water demand management strategies relies on our understanding of water consumers' behavior. This means inferring their consumption habits and routines, as well as identifying the most relevant determinants of water consuming or saving behaviors at the household level. Such information allow characterising the heterogeneity of water users' behaviors in a community, as well as building mathematical models that predict urban water demand variability in space and time, and explore the effects of different WDMSs for the residential sector.

In this chapter, we first contribute a novel descriptive modeling procedure, based on a combination of clustering and principal component analysis, which allows performing water users' segmentation on the basis of their eigenbehaviors (i.e., routines, recurrent water consumption behaviors) automatically iden-

¹The content of this chapter has been adapted from the following papers:

- Cominola, A., Moro, L., Riva, A., Giuliani, M., Castelletti, A., 11 – 14 July 2016b. Profiling residential water users' routines by eigenbehavior modelling. In: Proceedings of the 8th International Congress on Environmental Modeling and Software (iEMSs 2016). Toulouse (France);
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., Anda, M., 28 June – 3 July 2015b. Modelling residential water consumers' behaviors by feature selection and feature weighting. In: Proceedings of the 36th IAHR World Congress. The Hague (the Netherlands).

tified from smart metered consumption data. The approach is tested against a dataset of smart metered water consumption data from 175 households in the municipality of Tegna (CH). Numerical results demonstrate the potential of the method for identifying typical profiles of water consumption, which allow clustering water consumers' in heterogeneous groups.

Secondly, we also contribute a predictive modeling approach, based on feature selection and feature weighting to model the single-user consumption behavior at the household level. A two-step procedure consisting of the extraction of the most relevant determinants of users' consumption and the identification of a predictive model of water consumers' profile is proposed and tested on a real case study with low resolution data of over 1500 water users from Western Australia. Results show the effectiveness of the proposed method in capturing the influence of candidate determinants on residential water consumption, as well as in attaining sufficiently accurate predictions of users' consumption profiles, which constitutes essential information to support residential water demand management.

4.1 Introduction

The third phase of the general procedure to study residential water demand management we proposed in Chapter 2 (Figure 2.2, and also reported in Cominola et al. (2015a)) consists in modeling water consumers' behaviors. The goal of this phase is characterizing water demand at the household level, thus preserving the heterogeneity of the individual users in the modelled community. Single-household water consumption data provide essential information for accurately modeling individual users' behaviors, especially through the application of data analytics and machine learning techniques (e.g., Cardell-Oliver, 2013a). Two distinctive approaches exist for modeling water users' consumption behaviors: *descriptive models*, which aim at performing users' segmentation through the analysis of observed water consumption patterns and historical trends (e.g., Beal et al., 2011b; Beal and Stewart, 2014), possibly at the end-use level, and *predictive models*, which instead provide estimates of the expected water demands, possibly conditioned upon natural and socio- psychographic factors, or in response to alternative water demand management strategies (e.g., Maggioni, 2015; Makki et al., 2015).

The first class of models allows building users' consumption profiles based on historical trends, thus without requiring big database of information characterizing users features in terms of demography, economy, education, social habits, and attitudes. This provides the baseline reference for identifying promising areas where water savings and conservation actions may be focused. Yet, these models do not quantify the expected impact of demand management ac-

tions on water consumption and savings. In contrast, the second class of models can be employed to effectively predict water consumption at the household level, but at the same time require big amounts of users' features for model building purposes. Many state-of-the-art studies (e.g., Olmstead et al., 2007; Fox et al., 2009; Makki et al., 2013) reported the presence of correlations between one or more presumed consumption drivers and the associated consumption profiles, thus accomplishing the user profiling phase. Yet, the number of considered candidate variables is generally limited. In addition, only a few works completed a second phase of behavioral modeling and provide a quantitative prediction of the water demand at the household level as a function of the identified drivers and WDMSs, thus representing promising decision-aiding tools for water utilities and urban planners.

In this chapter, we propose a double contribution. First (see Section 4.2), we contribute a descriptive modeling procedure for performing users' segmentation on the basis of smart metered water consumption data. Our procedure is based on a combination of clustering and principal components extracted from water demand data, extending the idea originally proposed by Eagle and Pentland (2009) for the identification of routines in the temporal location of 100 individuals from MIT, monitored using 100 Nokia 6600 smart phones. The extraction of principal components from behavioral data defines a set of vectors spanning the "behavioral space" of monitored individuals, characterizing their behavioral variation in time. These components, called *eigenbehaviors*, are computed as the eigenvectors of the covariance matrix of behavior data, where the vectors associated to high weights represent a type of recurrent behavior, i.e., a *routine*. In this work, this idea is extended for the identification of typical water consumption behaviors from a dataset of smart metered water consumption readings. The proposed approach is tested on a dataset of hourly-sampled water consumption records from 175 households in the municipality of Tegna (CH), which have been equipped with smart meters by Società Elettrica Sopracenerina as part of the Smarth2O Project.²

Secondly (see Section 4.3), we contribute a predictive approach based on feature extraction techniques (Guyon and Elisseeff, 2003) to model the single-user consumption behavior at the household level. The approach is based on a two-step procedure: (i) identification of the most relevant determinants of users' consumption profiles and (ii) predictive model building of water consumption profiles based on the observation of the determinants identified in the previous step. The use of feature selection (i.e., algorithms returning a subset of selected features) and feature weighting (i.e., algorithms ranking the features according to their relevance) (Zhao et al., 2010) is motivated by the need to manage a large number of potentially relevant factors influencing water consumers' behaviors

²Smarth2O official website: www.smarth2o-fp7.eu.

along with their redundancy and highly nonlinear relationships, which represent major challenges for standard cross-correlation analyses (Galelli et al., 2014). The proposed approach is applied to a dataset of low-resolution water consumption records associated with a variety of demographic and psychographic users data and household attributes collected in nine towns of the Pilbara and Kimberley Regions of Western Australia throughout the H2ome Smart project (Anda et al., 2013).

4.2 Profiling residential water users' routines by eigen-behavior modelling

4.2.1 Methodology

The descriptive modelling procedure we propose in this section is composed of three main methodological phases:

- Data pre-processing;
- Eigenbehavior extraction through Principal Component Analysis;
- Eigenbehavior clustering through k-means.

Each phase is detailed in the next subsections.

4.2.1.1 Data pre-processing

Let's consider the case we have, for each smart metered household, a time series of water consumption readings sampled with hourly resolution, which can be organized in a $[D \times 24]$ individual water consumption matrix C^k (with $k = 1, \dots, U$, being U the total number of users/households), where each row corresponds to one day and each column to one hour of the day.³ We then transformed this matrix into a binary matrix Γ^k , where the observed values of hourly water consumption are classified into N mutually exclusive classes Λ (e.g., low, medium, high consumption) based on hourly consumption thresholds. Each row of Γ^k hence contains $24 * N$ elements, where the binary values in the i -th row of Γ^k associate the 24 consumption readings of the i -th day in C^k to one of the N consumption classes Λ . The generic element x of matrix Γ^k is hence defined as:

$$x(i, j + (n - 1) \times 24) = \begin{cases} 1, & C^k(i, j) \in \Lambda_n \\ 0, & \text{otherwise} \end{cases} \quad (4.1)$$

³We only consider total hourly water household consumption data, both in the methodological formalization and in the application, because at the time of this study no groundtruth for end-use algorithm calibration and appliance signature extraction are available. However, the methods we propose in this chapter can be easily extended to applications with availability of end-use data.

4.2. Profiling residential water users' routines by eigenbehavior modelling

where $i = 1, \dots, D$; $j = 1, \dots, 24$; $n = 1, \dots, N$. For each user, we can identify its average individual behavior Ψ^k as well as the daily deviation $\Phi^k(i)$ from this average behavior, defined as

$$\Psi^k = \frac{1}{D} \sum_{i=1}^D \Gamma^k(i) \quad (4.2)$$

$$\Phi^k(i) = \Gamma^k(i) - \Psi^k$$

4.2.1.2 Eigenbehavior extraction through Principal Component Analysis

After the data processing phase is completed, we can extract the eigenbehaviors by performing a Principal Component Analysis (PCA) (Jolliffe, 2002) on the resulting matrices Φ^k . PCA is a dimensionality reduction technique, which searches for linear combinations of the original variables such that the coefficients of the output combinations (the principal vectors) form a low-dimensional sub-space defined by directions explaining maximal variance in the original data. Few principal components usually explain a high percentage of the variance of the original variables, ensuring dimension reduction. In addition, the representation of the original data in the projected space defined by principal components is uncorrelated, thus providing a useful tool for physical and statistical interpretations.

PCA is performed via an eigenvalue decomposition of the covariance matrix R of Φ^k (which has D rows corresponding to D days and $Q = 24 * N$ columns corresponding to the binary labels classifying the hourly water consumption), i.e.

$$R = \frac{1}{D} \sum_{i=1}^D \Phi^k(i)^T \cdot \Phi^k(i) \quad (4.3)$$

where the resulting eigenvectors \mathbf{w}_q^k (with $q = 1, \dots, Q$) are the eigenbehaviors of the k -th user and allow mapping the original matrix Φ^k into its principal components, i.e.

$$\rho_q^k = \Phi^k \cdot \mathbf{w}_q^k \quad \text{with } q = 1, \dots, Q \quad (4.4)$$

4.2.1.3 Eigenbehavior clustering through k-means

After performing the dimensionality reduction step via PCA, we classify the metered users into different consumption profiles on the basis of their first eigenvector \mathbf{w}_1^k (with $k = 1, \dots, U$). The first eigenvector accounts for as much

4. Descriptive and predictive modelling of residential water users

as possible of the behavioral data variance of the considered user's consumption, which is quantified by the associated highest eigenvalue. Our classification was run using the K-means clustering method (MacQueen et al., 1967), a widely adopted technique that allows grouping multidimensional points in clusters by minimizing the average squared distance between points and centroid of each cluster. Formally, K-means algorithm partitions the U -dimensional set $\mathcal{W} = \{\mathbf{w}_1^1, \dots, \mathbf{w}_1^U\}$, containing the first eigenbehavior of each user, into $M < U$ profiles $\mathcal{P} = \{P_1, \dots, P_M\}$ by solving the following minimization problem:

$$\mathcal{P} = \arg \min_{\mathcal{P}} \sum_{m=1}^M \sum_{\mathbf{w}_1^k \in P_m} \|\mathbf{w}_1^k - \mu_m\|^2 \quad (4.5)$$

where μ_m is the mean of $\mathbf{w}_1^k \in P_m$, with the resulting clusters satisfying the following conditions: (i) the union of all clusters contains all the original point, i.e. $\cup_{m=1}^M P_m = \mathcal{W}$; (ii) each point belongs to a single cluster, i.e. $P_i \cap P_j = \emptyset$; (iii) the clusters cannot be empty and a single cluster cannot include all the points, i.e. $\emptyset \subset P_m \subset \mathcal{W} \quad \forall m$.

4.2.2 Case study

Our testing case study consists of real-world dataset of anonymized smart metered water consumption provided by Società Elettrica Sopracenerina (SES) as part of the SmartH2O project (see Rizzoli et al., 2014, for an overview about the project). SES, a multi-utility based in Locarno (CH) installed 400 smart meters in the Locarno district during the first two years (2014-15) of the SmartH2O project.

For this study, we consider a dataset that, after some pre-processing needed for removing missing readings, measuring errors, etc., comprises 175 households monitored for around 7 months, from March, 30th to November, 7th, 2015. Each user is therefore associated to 5,352 hourly readings (223 days), with the full dataset including 936,600 data.

4.2.3 Numerical results

The first step of our methodology required to transform the hourly water consumption readings matrix of each user C^k (with $k = 1, \dots, 175$) into binary matrices Γ^k , where the consumption data are classified into mutually exclusive consumption classes (see Section 4.2.1.1). We initially observed that 46% of the data in our dataset are equal to 0. This can be easily explained by the fact that generally there is no consumption during night or when people are not at home (e.g., during working hours). Considering the distribution of the nonzero values (see Figure 4.1), we decided to partition these latter in three classes with respect

4.2. Profiling residential water users' routines by eigenbehavior modelling

to a threshold of 12 L/h, representative of low consumption levels (e.g., faucet, toilet), and a threshold of 100 L/h, which allows distinguishing medium consumption events (e.g., a 10-minute shower, efficient clothes washer programs) from high consumption ones (e.g., outdoor uses, inefficient devices). In summary, we classified the available hourly water consumption readings into the following four consumption classes, whose sizes are illustrated in Figure 4.2:

1. **L0**: hourly consumption $\chi = 0$ L/h;
2. **L1**: hourly consumption $\chi \in (0, 12]$ L/h;
3. **L2**: hourly consumption $\chi \in (12, 100]$ L/h;
4. **L3**: hourly consumption $\chi > 100$ L/h.




Figure 4.1: Empirical distribution of nonzero water consumption readings, with the shading representing the three selected classes L1, L2, L3.

According to the thresholds set for the above classification, the hourly water consumption matrices C^k are transformed into the corresponding binary matrices Γ^k , from which the eigenbehaviors of each user can be extracted via PCA. In this practical application we adopted non-centered PCA, keeping into account that the Γ^k are binary, thus all their columns have the same lower and upper bound. Results show that considering solely the first principal component for representing the behaviors of the 175 considered users allows explaining, on average, 50% of the variance in the original data. The second principal component contributes an additional 8%, while 10 principal components would be




Figure 4.2: Partitions of hourly consumption reading data across the four consumption classes L0, L1, L2, L3.

necessary to explain 75% of the variance, with a saturation to around 100% obtained with 60 principal components. The large share of variance explained by the first eigenbehavior, along with the significant difference between the contribution of the first and second component, confirms that working on the first eigenbehavior is sufficient for capturing a large share of the variability in users' behaviors, thus supporting effective users' segmentation.

As a last step of our methodology, we applied K-means clustering on the set of first eigenbehavior of each user, where high values identify recurrent and relevant behaviors (e.g., frequent concentration of consumption events of a given class in specific hours, including periods with no consumption). K-means clustering would allow classifying users with similar consumption routines, represented by similar values of \mathbf{w}_1^k , in the same cluster, supporting effective users' segmentation. To demonstrate the validity of our procedure, we report the results obtained setting the number of clusters equal to three after silhouette analysis, thus simply distinguishing low, medium, and high consumption profiles (additional experiments with higher number of clusters are reported in Moro and Riva (2016)). Results are illustrated in Figure 4.3, where each column represents the first eigenbehavior of a user and each row corresponds to the hour of the day for each consumption class. The users' classified within profile P1 are characterized by high values in the first eigenbehavior only for L0, meaning they are usually not consuming: this may be the case of vacation houses or, when the eigenbehavior coefficient is positive in classes L1 or L2, they are houses occupied by just one inhabitant, with low consumption. Also profile P2 shows high values for L0, but mostly during the night. In addition, this profile exhibits a bimodal behavior for L3 and non-null coefficients for L2. This identifies a typical situation of average consumers that do not consume in the

4.2. Profiling residential water users' routines by eigenbehavior modelling

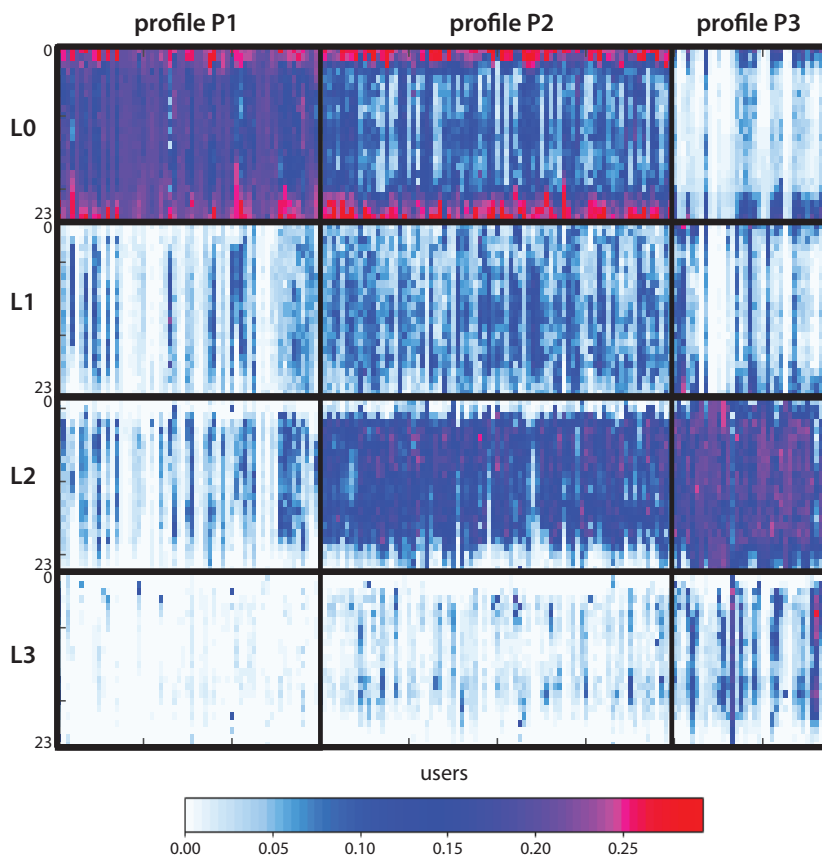


Figure 4.3: Results of users' segmentation with respect to the values of the first eigenbehaviors. Each column is the first eigenbehavior of a user and each row corresponds to the hour of the day for each consumption class L0, L1, L2, L3. Color scale represents the value of eigenbehaviors coefficients.

night, but have two peaks of medium water consumption in the morning and in the evening and low consumption (e.g., toilet) equally weighted during day hours. Finally, profile P3 is characterized by high values in the first eigenbehavior mainly associated to L2 and L3, again with a double peak in the morning and in the evening. This behavior might simply suggest houses in this cluster have larger size (in terms of number of occupants) than those in cluster P2. The main characteristics of each profile are illustrated in Figure 4.4 in terms of relative frequency of average hourly consumption for each profile in each consumption class. This analysis confirms that profile P1 is associated to very low consumers, with high frequency of no-consumption events. Profile P2 captures middle consumers that concentrate their water consumption in the morning and evening (see the double peak in the frequency of L2), with almost no con-

4. Descriptive and predictive modelling of residential water users

sumption during night and few events in class L3. Finally, profile P3 identifies high consumers, characterized by a double peak of high consumption (class L3) and frequent consumption events of class L2 during the entire day.

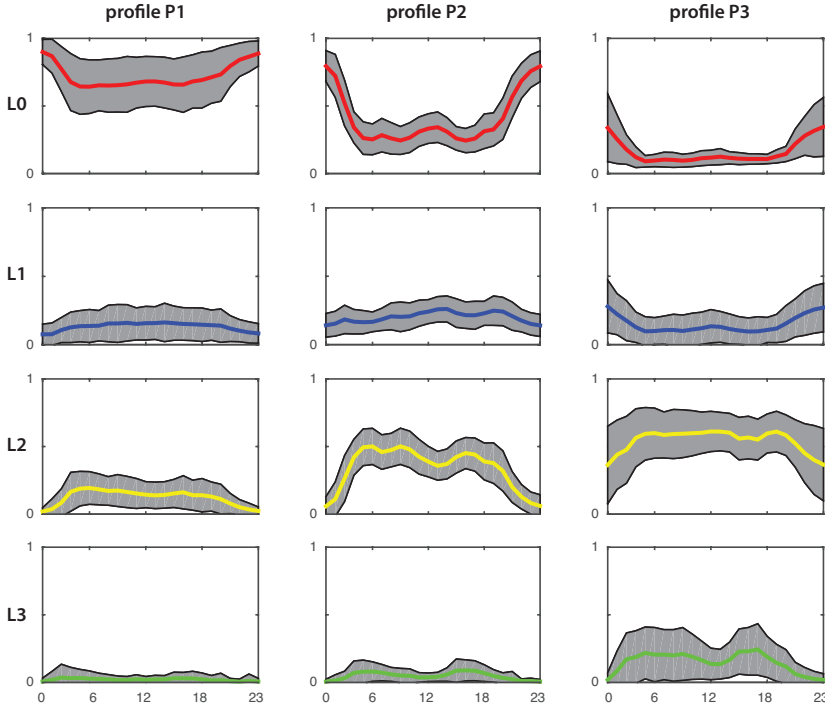


Figure 4.4: *Relative frequency of average hourly consumption for each profile in each consumption class. The gray area represents standard deviation.*

4.3 Modelling residential water consumers' behaviors by feature selection and feature weighting

4.3.1 Methodology

Feature extraction techniques, mostly developed in the data mining and machine learning research communities, represent potentially promising tools to model residential water users behaviors. These methods allow extracting the more relevant determinants in describing the consumption profiles of water users out of a large set of candidate drivers. On the basis of the selected determinants, a behavioral model predicting the water consumption at the household level can be identified. The general formulation of a water demand predictive model for a generic user i is the following:

$$\hat{y}_i = f(x_i) \quad (4.6)$$

where \hat{y}_i is the consumption profile of the i -th user and x_i denotes the set of M determinants influencing his/her behavior, represented by a variety of demographic and psychographic users data (e.g., age, number of house occupants, income level, conservation attitude, etc.), household attributes (e.g., house size, type, garden area, etc.) and exogenous factors (e.g., temperature, and precipitation, water price, etc.). The union of determinants and consumption data yields a sample dataset containing N tuples, one for each user. The i -th tuple (with $i = 1, \dots, N$) is defined as follows:

$$\langle x_i^1, x_i^2, \dots, x_i^M, y_i \rangle \quad (4.7)$$

The construction of the water demand predictive model defined in (4.6) relies on the following two-step procedure:

1. Feature extraction, to select from the original dataset X of users' data a subset $X' \subseteq X$ of determinants that are relevant to describe the consumption profile Y ;
2. Model learning, to predict the water consumption profile as a function of the selected features X' .

4.3.1.1 Feature extraction

Different approaches can be adopted to perform feature extraction as well as for model learning. In particular, feature extraction techniques can be classified in two main categories:

- *FS*, namely algorithms that return a subset of features selected from the original dataset as the most relevant to describe the considered output variable (i.e., consumption profile);
- *FW*, namely algorithms that rank all the features according to a measure of their relevance, with no actual selection of the most relevant variables, which however are identified as the ones in the first positions of the ranking.

Moreover, depending on their structure, they can be distinguished between *model-free* (or *filter*) algorithms, when they do not include any learning algorithm, or *model-based* in case they explicitly rely on a learning algorithm (Galelli and Castelletti, 2013). Model-based algorithms can be further classified into wrapper models, if they include a predetermined learning algorithm, and embedded, if the model construction phase includes feature selection (Zhao et al., 2010). Since no single method is best suited to all datasets and modeling purposes a-priori, we implemented and applied different algorithms for both feature selection and weighting.

Feature selection algorithms We run the following feature selection algorithms:⁴

- Fast Correlation Based Filter (FCBF) (Yu and Liu, 2003). This algorithm exploits the formulation of the Information Gain algorithm (explained in the next section among the feature weighting algorithms), in order to keep into account both the feature-feature and the feature-class correlation, thus considering redundancy issues. The correlation between a feature X_i and a class C is computed through the concept of Symmetrical Uncertainty (SU), which is defined as follows:

$$SU(X_i, C) = 2 \frac{IG(X_i, C)}{H(X_i) + H(C)} \quad (4.8)$$

where IG represents the information gain and H is the entropy of a variable (as defined later on).

- Correlation Feature Selection (CFS) (Zhao et al., 2010). This filter algorithm uses a correlation-based heuristic to determine the relevance of a feature both in terms of feature-class correlation and feature-feature intercorrelation, thus avoiding redundancy issues. Given a subset of n features, the algorithm determines the "worth of the subset" and then explores different subsets in order to identify the one with the best merit.
- Bayesian Logistic Regression (BLOGREG) embedded method (Guyon et al., 2002). This an embedded feature selection algorithm, which promotes the sparsity of a logistic model in order to reduce the number of features selected. The BLOGREG algorithm is suitable for managing categorical features.
- Sparse Bayesian Multinomial Logistic Regression (SBMLR) embedded method (Cawley et al., 2007). This algorithm is an extension of the BLOGREG approach, making it suitable to be applied to multiclass problems.

We also tested the following feature weighting algorithms:

- CHI-square score (Liu and Setiono, 1995). This is a supervised, filter algorithm used to test whether the class labels are independent of a specific feature. In particular, the chi-square score is evaluated as follows:

$$\chi^2 = \sum_{i=1}^k \sum_{j=1}^{N_{\text{class}}} \frac{(n_{ij} - \mu_{ij})^2}{\mu_{ij}} \quad (4.9)$$

where:

⁴The 2014 version of the ASU feature selection package downloadable at <http://featureselection.asu.edu/> was adopted for this study.

- n_{ij} is the number of samples with the i -th value for a particular feature in class j ;
- $\mu_{ij} = \frac{n_{*j}n_{i*}}{n}$, with n_{*j} being the number of elements with value i for a particular feature across all classes and n_{i*} being the number of elements in class j . The higher chi-square score, the more the class label is dependent on the considered features. The index is suitable to be applied also with categorical (or binary) variables.
- Information gain (Cover and Thomas, 2012). Information gain is another measure of dependence between a feature and the class labels. Considering a feature X_i and the class labels C , the information gain is defined as:

$$IG(X_i, C) = H(X_i) + H(X_i|C) \quad (4.10)$$

Where H represents the entropy of a variable, defined as:

$$\begin{aligned} H(X_i) &= - \sum_j P(x_j) \log_2(P(x_j)) \\ H(X_i|C) &= - \sum_z P(y_z) \sum_j P(x_j|y_z) \log_2(P(x_j|y_z)) \end{aligned} \quad (4.11)$$

It is worth mentioning that both the chi-square score and the information gain algorithms consider each feature separately, thus they do not solve redundancy issues.

4.3.1.2 Model learning

As far as the model learning is concerned, in principle any data-driven model (regressors or classifiers) can be used see (see Maier and Dandy, 2000; Galelli and Castelletti, 2013). In practice, the selected method should have the following desirable features: (i) *modeling flexibility* to approximate strongly non-linear functions, particularly because the relationships between the candidate inputs (selected features) and the output (consumption profile) is completely unknown a priori; (ii) *computational efficiency* to deal with potentially large data-sets, when considering large number of users; (iii) *scalability* with respect to the number of candidate variables to be analyzed, due to the need of testing several variables with different domains and variability. In the present experiments, we used two different data-driven models:

- a Naive Bayes Classifier (Duda et al., 1973). Bayesian classifiers learn from training data the conditional probability of each attribute, given the class label. Classification is then performed by computing the probability of each class, given an instance of the attributes and predicting the class with the highest posterior probability;

4. Descriptive and predictive modelling of residential water users

- the J48 java implementation of the C4.5 Decision Tree algorithm (Quinlan, 2014). The J48 algorithm used here is an implementation of the C4.5 algorithm used to generate a decision tree. It builds a decision tree on the training dataset, where the attributes that most effectively split the set of samples into small subsets, in terms of information gain, are positioned onto nodes.

4.3.2 Case study

4.3.2.1 The H2ome Smart project

The following data, collected within the H2ome Smart project, are available:

- **Household water consumption data** from meter readings (measured in m^3). The maximum number of readings per household within the considered period is seven, thus the highest reading resolution is approximately three months;
- **House and occupants attributes:** 26 variables describing different features of the users and their house. Table 4.1 reports the complete list of data available.

Data were collected between August 2010 and February 2012 for more than 3000 households in the towns of the Pilbara and Kimberley Regions of Western Australia.

4.3.2.2 Data pre-processing

Data cleaning We performed the following data cleaning operations:

1. Records of users showing data inconsistencies or missing data (i.e., negative consumption rate or no consumption rate measures for any reading period) were removed from the dataset;
2. Empty reading dates fields were filled for as many users as possible with the reading dates of the same accounting reading group;
3. The average daily water consumption rate in $[\text{m}^3/\text{day}]$ was computed for each water-reading period, for each household, given its water consumption data and reading period dates. This operation was useful to obtain comparable values of water consumption among different houses, since the duration of reading period was heterogeneous in the considered sample;
4. If information about the number of house occupants was present, the per-capita daily water consumption rate in $[\text{m}^3/\text{day}]$ was computed for each reading period.

4.3. Modelling residential water consumers' behaviors by feature selection and feature weighting

Table 4.1: *Water consumers' and household features list.*

Variable name	Description	Variable nature	Number of categories
Town	-	categorical	9
Suburb	-	categorical	21
Years of occupancy	Years since the house is being occupied by the same inhabitants	integer	-
House responsibility	Person responsible for paying bills	categorical	4
Number of occupants	Number of inhabitants in the house	integer	-
Resident type	Type of resident in the house	categorical	8
Number of toilets	Number of toilets in the house	integer	-
Land use	Type of land use destination	categorical	14
House type	Type of house structure	categorical	5
Washing machine type	Type of washing machine	categorical	3
Toilet type	Type of flush	categorical	3
Shower type	Type of shower	categorical	3
Dishwasher presence	Presence of dishwasher	binary	-
Garden area	Area of the house garden [m ²]	real positive	-
Watering method	Method used for garden watering	categorical	4
Watering time	Average weekly watering time	integer	-
Irrigation system	Type of irrigation technique	categorical	3
Drip irrigation type	Type of irrigation technique	categorical	3
Surface irrigation type	Type of surface irrigation	categorical	3
Drip irrigation duration	Weekly average drip irrigation minutes	categorical	4
Surface irrigation duration	Weekly average surface irrigation minutes	categorical	4
Mulch usage	Usage of mulch	binary	-
Pool presence	Presence of pool	binary	-
Pool cover usage	Presence of pool cover	binary	-
Spa presence	Presence of spa	binary	-
Native plants presence	Presence of native plants	binary	-

4. Descriptive and predictive modelling of residential water users

The data cleaning process produced the following outputs: a matrix C_{daily} containing six readings of daily average water consumption rate for $N = 1624$ households and a matrix C_{pcDaily} containing six values of per-capita daily average consumption for $N' = 1560$ households. Note that N and N' are significantly lower than the initial dimension of the dataset, which included approximately 3000 households, as water consumption readings were partially or totally missing.

Class label assignment The real values in C_{daily} and C_{pcDaily} were converted into three classes representing different water consumption profiles: low-consumers, medium-consumer, and high-consumers. K-means clustering was used to assign consumption data to classes, with $k=3$ (number of classes) and Squared Euclidean distance settings. It was run over the vectors Y_{daily} and Y_{pcDaily} containing, respectively, the mean of water readings in C_{daily} and C_{pcDaily} , for each household.

Matrix of users' and households features Two sample datasets X_{daily} and X_{pcDaily} were built, respectively for the users whose consumption is included in Y_{daily} and Y_{pcDaily} . Each tuple of the datasets has $M = 26$ user and house features (see Table 4.1) associated to either Y_{daily} or Y_{pcDaily} . The processed datasets X_{daily} and X_{pcDaily} consisted, respectively, of N and N' tuples, one for each user satisfying the pre-processing conditions.

4.3.3 Numerical results

4.3.3.1 Feature selection and feature weighting

The outputs of the feature selection algorithms are represented in Figure 4.5, where the user and house features are listed on the y-axis and the color indicates the selection frequency of each feature over different algorithms runs.

Each feature extraction algorithm was run 3 times: both X_{daily} and X_{pcDaily} were split into three subsets of equivalent size, and each run considered two thirds of the dataset for feature extraction calibration and the remaining third for predictive modeling validation (Section 4.3.4). Dark colored features in Figure 4.5 are the most relevant as they are always selected across the different algorithms runs, while their relevance decreases moving towards gray and white tones. The results of the two figures appear to be quite consistent: the number of household's occupants seems to be the most relevant factor impacting average daily residential water consumption Y_{daily} (left part of the figure), as its selection frequency is higher than 80%; the number of toilets, the method used for irrigation, the presence of a pool and the land use destination are then ranked in the subsequent positions, with a decreasing selection frequency, but

4.3. Modelling residential water consumers' behaviors by feature selection and feature weighting

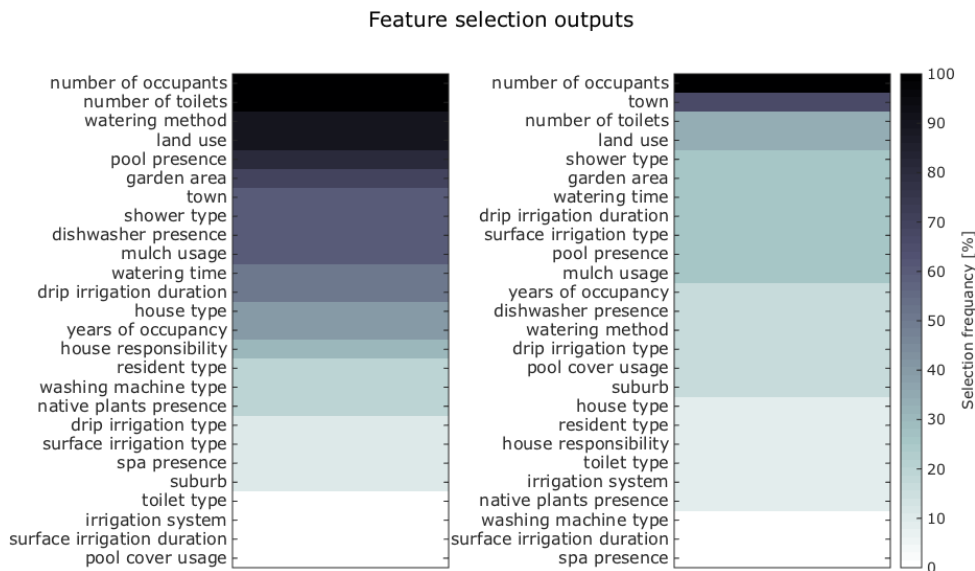


Figure 4.5: Selection frequency of candidate features over multiple feature selection algorithms runs. Considered predictant: Y_{daily} (left) and Y_{pcDaily} (right).

still higher than 60%. In addition, the geographical position, expressed by the "town" attribute, is also considered relevant in explaining the average per-capita daily water consumption Y_{psDaily} (right part of the figure). However, the results obtained considering the average per-capita daily water consumption as predictant enforce the relevancy of the number of house occupants as main driver of water consumption, since its selection frequency is 100%, while all the other candidate variables do not achieve a selection frequency higher than 70%.

Figure 4.6 shows the results obtained by running the feature weighting algorithms on X_{daily} and X_{pcDaily} , respectively. Again, the features are reported on the y-axis, while the x-axis represents the different feature weighting algorithms runs. Colors represent the positions of each feature in the weighting ranking: features with dark color were given higher weights by the algorithms, meaning they are considered relevant in explaining the output variable, while lighter features are associated to lower weights (i.e., less relevant). The two feature weighting algorithms produce consistent results, which are also coherent with the ones obtained by the feature selection algorithms, at least for the majority of the top-ranked features. The results confirm the existence of clear and

4. Descriptive and predictive modelling of residential water users

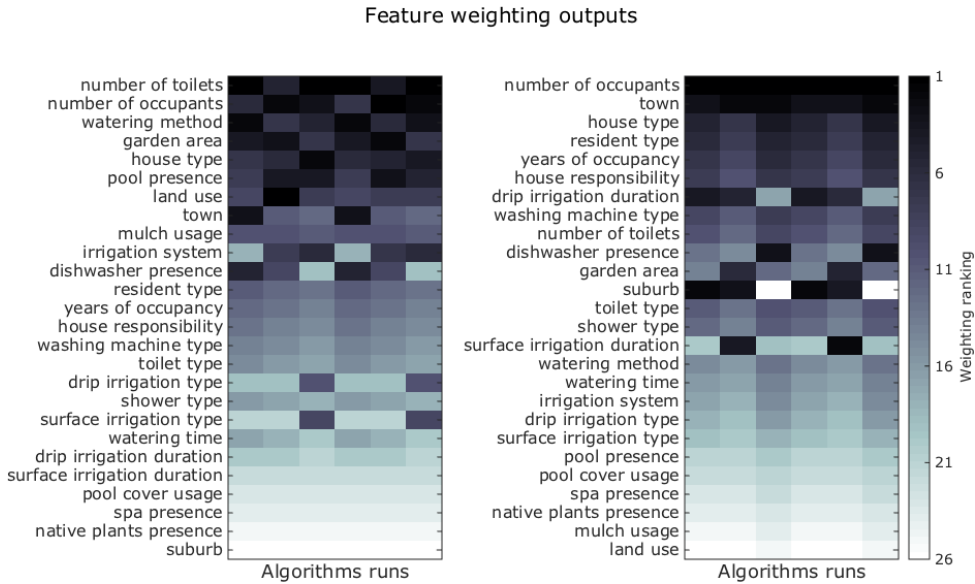


Figure 4.6: *Weighting ranking of candidate features over multiple feature selection algorithms runs. Considered predictant: Y_{daily} (left) and Y_{pcDaily} (right).*

strong relationships between the extracted features and the corresponding water consumption profiles.

4.3.3.2 Interpretation of the feature extraction results

In this section, a set of top-ranked features extracted in the previous section is analyzed, in order to better understand the underlying relationships between them and the water consumption profiles.

Number of occupants The first considered feature is the number of occupants of the house, which is always ranked in the first position by all the algorithms. As shown in Figure 4.7, not surprisingly the median daily water consumption increases with the number of occupants. Yet, the median per-capita consumption decreases with the increasing of the number of occupants. The reason for that can be twofold: the first reason might be that some end-uses represent a sort of fixed-cost, which is shared among the occupants. For example, the water used for irrigation or in a pool is shared among the occupants and, therefore, the individual cost (i.e., consumption) decreases for increasing

4.3. Modelling residential water consumers' behaviors by feature selection and feature weighting

number of inhabitants. The second reason might be that when the number of occupants increases, some kind of economies of scale and social pressure develop. As a consequence, water use is better balanced among the inhabitants and wastes are less frequent (Beal et al., 2011a).

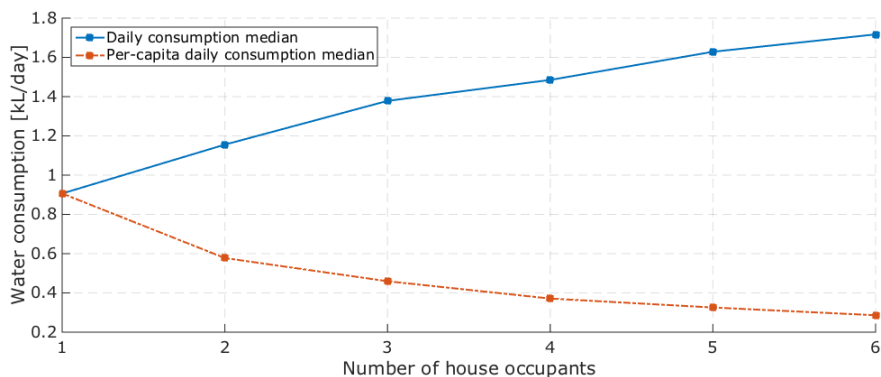


Figure 4.7: Median daily water consumption and median per-capita daily water consumption for different numbers of house occupants.

Number of toilets Figure 4.8 analyzes the number of toilets, where both the median daily and median daily per-capita water consumption level increase with the number of toilets in the house. Since the number of toilets generally increases with the size of the house (and thus with the number of household's occupants), it is reasonable that the daily water consumption increases with the number of toilets. In contrast with the previous case, also the median per-capita consumption increases, probably because with a higher number of toilets there is less "competition" for using the resource (i.e., the toilet).

Irrigation and garden size The relationship between water consumption and the type of irrigation is shown in Figure 4.9: households where irrigation is performed by hand consume (on average) less water than those houses where irrigation is performed with automatic irrigation systems or both by hand and automatically. This evidence can be explained by relating the water consumption levels to the area of the garden to be irrigated (right y-axis). Houses equipped with automatic irrigation systems generally have a wide garden and high water consumption for irrigation. On the contrary, small gardens are irrigated by hand, resulting in a lower consumption. Reasonably, in houses with a medium-size garden and medium consumption levels, irrigation can be either manual or automatic.

4. Descriptive and predictive modelling of residential water users

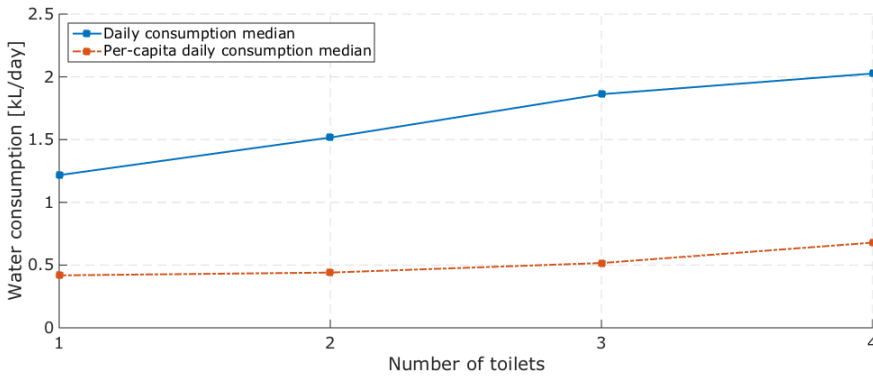


Figure 4.8: Median daily water consumption and median per-capita daily water consumption for houses with different number of toilets.

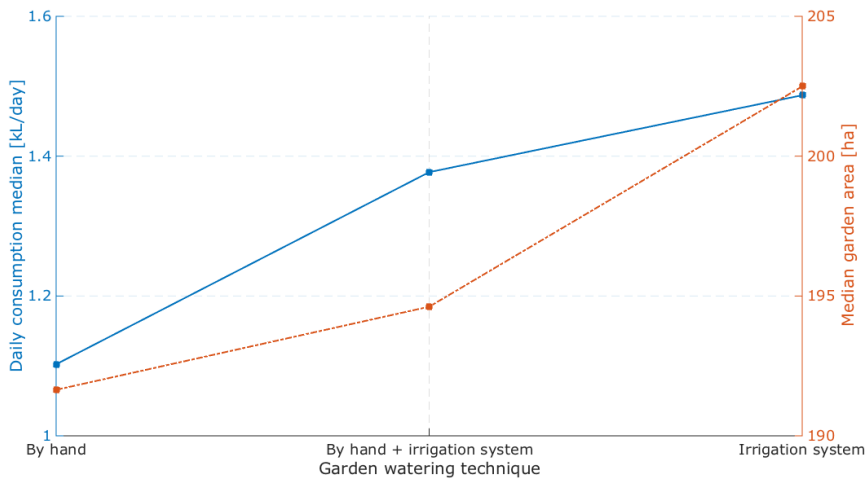


Figure 4.9: Median daily water consumption and median garden area for households adopting different irrigation techniques.

Type of house Figure 4.10 shows how the consumption level increases with the size of the house. This phenomenon can be probably explained as bigger houses generally are occupied by a higher number of inhabitants and, also, they have a higher number of toilets or very likely larger gardens. In turn, the per-capita water consumption flattens for the same reasons previously discussed about the relationship between the number of occupants and their associated per-capita consumption.

4.3. Modelling residential water consumers' behaviors by feature selection and feature weighting




Figure 4.10: Median daily water consumption and median per-capita daily water consumption for different types of house.

4.3.4 Model learning and users' water consumption profile forecasting

The second step of our procedure aims at identifying a model having the features extracted in the previous section as input, and the predicted water consumption profile of the users as output. This second step is fundamental in order to properly support the design of WDMSs, as well as assess their effectiveness. Indeed, the first step of feature extraction provides indications of potentially relevant water consumption drivers, thus supporting the empirical design of WDMS with a description of the status quo of users' behaviors. In turn, a predictive model able to forecast consumers' profile based on relevant attributes enables quantifying changes in household water consumption due to modifications in natural and socio-demo-psycho-graphic drivers, thus supporting utilities and planners by anticipating the effect of WDMSs. In particular, working on low-resolution consumption data, our model allows classifying users to the three consumption profiles introduced in Section 4.3.2.2, namely low-, medium-, high-consumers. Among the available data-driven models, we employed Naive Bayes Classifier and Decision Tree algorithm (see Section 4.3.1.2) which are particularly suitable for these classification experiments. In order to minimize the risk of overfitting the model over the calibration data, we run a k-fold cross-validation, with $k=3$, by randomly splitting the dataset into k mutually exclusive subsets of equivalent size. Each time the predictive model is validated on one of the k folds and calibrated using the remaining $k-1$ folds, on which the feature extraction algorithms are run. Table 4.2, Table 4.3 and Figure 4.11 report the resulting average models accuracy across the k-fold cross-validation, measured in terms of percentage of correct assignments of users on the basis of their features to their actual consumption profile.

4. Descriptive and predictive modelling of residential water users

Table 4.2: *Naive Bayes Classifier prediction accuracy based on FS algorithms outputs.*

FS algorithm	Average Naive Bayes Classifier accuracy on Y_{daily} [%]	Average Naive Bayes Classifier accuracy on Y_{pcDaily} [%]
FCBF	62.11 ± 2.80	80.45 ± 7.75
CFS	63.22 ± 2.91	80.45 ± 7.75
BLogReg	62.48 ± 3.83	80.45 ± 7.75
SBMLR	63.03 ± 3.11	80.45 ± 7.75

Table 4.3: *J48 Decision Tree prediction accuracy based on FS algorithms outputs.*

FS algorithm	Average J48 Decision Tree accuracy on Y_{daily} [%]	Average J48 Decision Tree accuracy on Y_{pcDaily} [%]
FCBF	63.46 ± 1.99	80.64 ± 7.89
CFS	64.94 ± 2.71	80.64 ± 7.89
BLogReg	61.92 ± 0.85	80.00 ± 7.35
SBMLR	62.91 ± 1.72	81.15 ± 7.84

4.3. Modelling residential water consumers' behaviors by feature selection and feature weighting

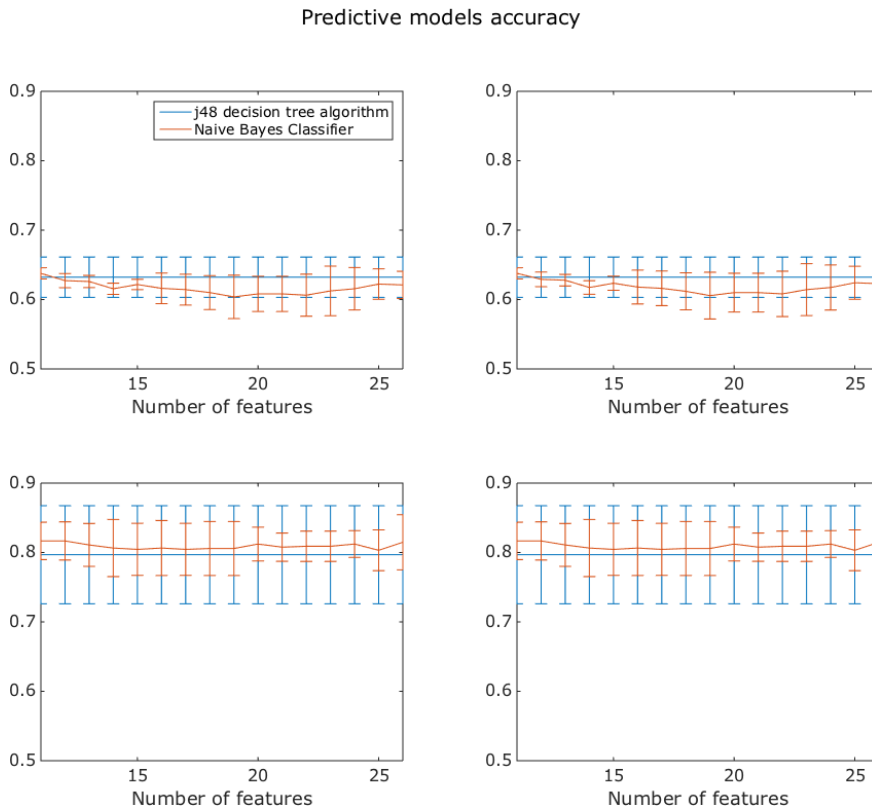


Figure 4.11: Predictive models accuracy based on FW algorithms outputs. The following FW algorithm predictant are represented: Y_{daily} and Information gain FW (top-left), Y_{pcDaily} and Information gain FW (top-right), Y_{daily} and Chi-square score FW (bottom-left), Y_{pcDaily} and Chi-square score FW (bottom-right).

Results show that both the models allow attaining a sufficiently good accuracy in predicting the consumption profiles of the users, both when users are classified according to the total consumption of their house or the per-capita consumption level. Moreover, although Figure 4.11 shows that the prediction accuracy slightly varies when the number of features considered in the model increases, feature extraction algorithms succeeded in identifying the smallest subset of most relevant features, allowing for a sufficient level of prediction accuracy. The proposed method hence shows the potential to effectively capture urban water demand variability with respect to users psychographics and house characteristics data, thus representing promising decision-aiding tools for water utilities and urban planners.

4.4 Discussion

In this chapter we presented a descriptive and a predictive modelling tool, in order to support the design of WDMSs through the understanding of single-user consumption behavior at the household level. First (see Section 4.2), we proposed a novel procedure for performing residential water users' segmentation from smart metered consumption data. Our procedure is based on a combination of clustering and principal component analysis and allows the automatic identification of recurrent water consumption behaviors in the form of eigen-behaviors. The approach was tested on a dataset of smart metered water consumption data from 175 households in the municipality of Tegna (CH). In the second part of the chapter (see Section 4.3), we proposed a two-step procedure consisting of the extraction of the most relevant determinants of users' consumption profiles, through feature extraction techniques, and the identification of a predictive model of water consumers' profile. We tested the procedure against a dataset containing low-resolution water consumption records associated with a variety of demographic and psychographic users' data collected within the H2ome Smart project, in nine towns located in Western Australia.

Numerical results show that the descriptive approach successfully extracts the main routines characterizing the metered users. Indeed, the heterogeneous profiles identified, even from a small sample composed of less than 200 users, represent typical consumption patterns (Cardell-Oliver et al., 2016) reflecting different behavioral habits of the users. This segmentation seems promising for informing the design of customized water demand management strategies. As for the predictive approach, results show overall consistency among the feature extraction techniques applied. The analysis of the results allows understanding the relationships between the selected features and the consumption profiles, demonstrating the suitability of such techniques as tools for capturing the influence of candidate determinants on residential water consumption. The development of predictive models of users' behavior attains sufficiently high accuracy in predicting the household water consumption as a function of the user features, which constitutes essential information to support residential water demand management strategies.

Further analysis will focus on comparing the proposed methodologies with other user profiling techniques, as well as assessing the robustness of our results and testing the sensitivity of each modelling components, for both approaches. For instance, preliminary tests show that the clustering technique used for the construction of the users' consumption profiles impacts on the final results of the predictive model. Moreover, we will assess how the overall descriptive and predictive accuracies might vary when moving from low-resolution billed data/high-resolution total consumption to high-resolution smart-metered data disaggregated for each end-use, which would allow the definition of more de-

tailed user profiles on the basis of disaggregated end-use patterns. The entire user profiling process would then benefit from the use of alternative methods for direct interaction with the users for data gathering as well as for providing personalized feedbacks, as the psychographic users data and the house characteristics considered in this chapter were collected via survey with no guarantees that all the relevant determinants of users' behaviors are observed and updated along time.

Finally, the integration of the descriptive and the predictive capabilities of the models we proposed here and applications of the integrated procedure on larger datasets, possibly involving hundreds or thousands water users from different contexts, and including an exhaustive set of psychographic features, as well as end-use energy and gas consumption data, would definitely contribute a significant step-forward in the extraction of relevant information for decision-making from big data of resources use in the residential sector.

5

Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

Abstract¹

Customized DSM strategies complementing supply-side management to meet future water and energy demands, and reduce utilities' costs, can be designed for groups of heterogeneous users characterized by different demand intensity, time-of-use, and opportunities for conservation or more efficient uses. Proper data mining and modelling tools are essential to extract these characterizing information out of smart-metered data comprising several thousands of data points. In this chapter, we contribute a customer segmentation analysis of over 1000 residential accounts in the South California, with the goals of exploring heterogeneity of typical residential water-electricity demand profiles,

¹The content of this chapter has been adapted from the following paper in preparation: Cominola, A., Spang, E., Giuliani, M., Castelletti, A., Loge, F., Lund, J., 2016c. Residential water and electricity customer segmentation analysis to inform demand-side management programs. *Environmental Science & Technology* (in preparation).

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

interpreting them in terms of users' behavior, and providing insights to coordinated water-energy DSM. Our procedure combines eigenbehavior extraction for the identification of typical consumers profiles from smart metered hourly water-energy time series, clustering for customer segmentation based on profile similarities, and factor mapping to infer the potential determinants of targeted profiles to inform the design of customized DSM. Results show that daily water and electricity consumption magnitudes are correlated, yet they show different daily load shapes. This has key policy implications because it is suggesting that groups of high consumers can be targeted with coordinated water-electricity DSM interventions, e.g., dynamic pricing and block tariffs. Yet, results also show that there is not a relevant causal nexus between water and electricity demands: this is suggesting that DSM actions for water should be different than those for electricity, even though target user groups are common. Finally, we found that both objective (e.g., presence of swimming pool) and subjective psychographic factors (e.g., conservation attitude) are relevant potential drivers of water-electricity demands. Based on that, we propose recommendations for designing a portfolio of mixed customized water-electricity DSM to foster conservation or peak shifting objectives.

5.1 Introduction

In Chapter 2 we illustrated a general procedure to study residential water (and energy) demand management relying on smart-metered data (see Figure 2.2). After gathering high-resolution consumption data and modelling users' behavior based on such data, possibly disaggregated at the end-use level, we mentioned as a last phase of the process the design of personalized demand management strategies, to modify users' demand in order to make it more efficient. Indeed, DSM strategies are key to pursue several short-to-long term objectives, both in the water and energy field. More precisely, they are useful to complement supply-side management in order to meet future demands, pursue resources conservation (e.g., conserving water in areas affected by prolonged periods of resource scarcity such as California (Arbués et al., 2003; Cahill and Lund, 2012)), they are cost-effective in supporting utilities business (Gleick et al., 2003) and reduce their operational short and long term costs (e.g., through peak consumption reduction (Cardell-Oliver et al., 2016; Beal et al., 2016)), and foster water distribution and consumption efficiency.

In the previous chapters of this thesis, we proposed a number of modelling and data mining tools that we demonstrated allow extracting information useful to characterize individual water consumers, retrieving meaningful information out of large smart-metered datasets. Our studies and other state-of-the-art works demonstrated these tools to be suitable to discriminate among different

daily demand consumption patterns (Espinoza et al., 2005; Verdú et al., 2006), thus identifying heterogeneous recurring (Kwac et al., 2014) or irregular behaviors (Nizar and Dong, 2009). Indeed, literature presents several user modelling approaches to DSM based on the combination of smart metering and customer segmentation. These studies have been performed especially in the energy sector, with some early applications to water data exploiting customer segmentation to identify routines in water consumption patterns and potential links with salient household characteristics (e.g., Cardell-Oliver et al., 2016). Findings from these works prove that the heterogeneity of consumers' profiles and behaviors identified by means of customer segmentation is key to design customized DSM strategies and inform utilities on representative demand patterns among their consumers' population. Yet, while smart meters also fostered the development of several consumers' models in the literature on water and energy demand management, on the other hand extracting users' profiles out of big smart-metered datasets, so that they provide an accurate representation and coverage of the consumption routines in the considered users' community, and at the same time are concise enough to inform water utilities and take demand management decisions, is a main challenge. Hence, DSM benefits from tools that discriminate among heterogeneous daily water/energy consumption profiles, while filtering out the high noise and variance of smart metered data caused by irregular and scattered users behaviors.

In addition, the water-energy nexus in the residential sector (Spang and Loge, 2015; Escriva-Bou et al., 2015a; WWAP, 2014) and its implications towards resources conservation and efficiency (Abdallah and Rosenberg, 2014; Escriva-Bou et al., 2015b) is pushing towards synergies for the development of coordinated water-energy demand management actions. Yet, to our knowledge no studies performing customer segmentation on coupled hourly water and electricity data, thus informing multi-utilities of coordinated water-energy demand management, have been performed so far in the literature.

In this chapter, we integrate part of the components of the models described in the previous chapter and contribute a customer segmentation analysis of over 1000 residential accounts in South California, with the goals of exploring differences between typical residential water and electricity demand patterns, interpreting these typical demand patterns in terms of users' behavior, and providing useful insights to coordinated water-energy demand management programs based on resources conservation or demand peak shifting. Each account is described by coupled household water and electricity hourly consumption data, as well as psychographic variables, i.e., variables related to occupants' demography, households characteristics and water/energy-related attitudes and stated preferences. We propose a novel descriptive modeling procedure for customer segmentation based on hourly water and electricity data. Our procedure inte-

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

grates a combination of eigenbehavior extraction (see Section 4.2 in Chapter 4) for the identification of typical routines, i.e., consumers profiles, in smart metered water-energy time series, k-means clustering for customer segmentation based on profile similarities, and factor mapping (Friedman and Fisher, 1999) to infer the potential determinants of targeted water and energy demand profiles to inform customized demand management. We show that our procedure captures the heterogeneity of water and electricity consumption profiles among the considered accounts, highlighting also differences between daily time-of-use of water and electricity, and allowing for the characterization of accounts with similar profiles based on psychographic variables and behavioral factors. Finally, we design recommendations for water and electricity conservation and efficiency programs targeting those accounts characterized by profiles significantly contributing to the total and peak water-energy consumption overall the community of households monitored.

5.2 Methodology

We developed and propose here a three-phases methodology, represented in Figure 5.1, to identify representative (i.e., recurring) water and energy users' profiles, perform consumers' segmentation based on their profiles, and infer the potential determinants of targeted water and energy consumption profiles. The methodology is composed of the phases described in the following sections.

5.2.1 Eigenbehavior extraction

The goal of this first step is characterizing each account by identifying its most recurring coupled daily water and electricity demand profiles. Its required inputs are hourly, smart-metered, water and electricity consumption time series, and it outputs the most recurring daily water-and-electricity consumption profiles, for each account. Further developing the idea originally proposed by Eagle and Pentland (2009) for the identification of routines in the location along time of 100 individuals monitored using smartphones, which we extended in Chapter 4 for the extraction of typical water consumption behaviors from a dataset of smart metered water consumption readings², we extract recurring coupled water and electricity daily demand profiles from smart metered water and electricity datasets by means of non-centered PCA (Jolliffe, 2002). PCA performs data dimensionality reduction, while searching for linear combinations of the original variables (in our case water and electricity consumption for each hour of the day) such that the coefficients of the output combinations (the principal vectors) form a low-dimensional sub-space defined by directions

²See also Cominola et al. (2016b)

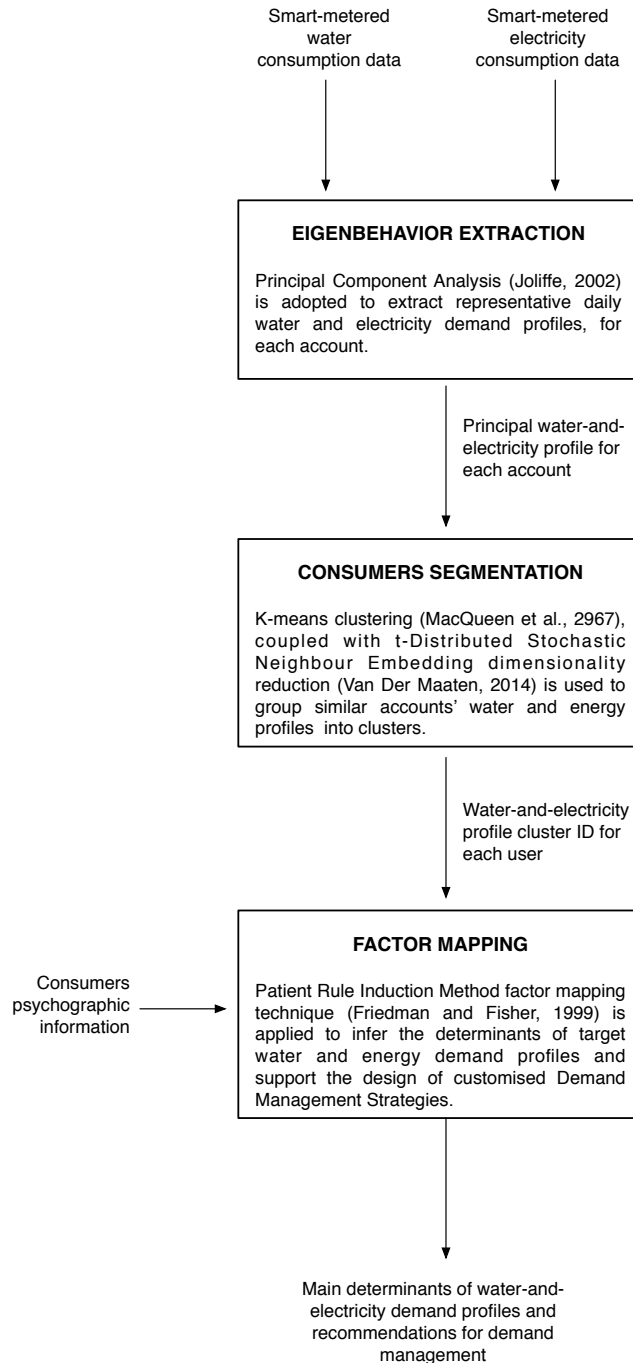


Figure 5.1: Flowchart of the methodology.

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

explaining maximal variance in the original data. The extraction of principal components from behavioral data defines a set of vectors spanning the "behavioral space" of monitored individuals, characterizing their behavioral variation in time. These components, called eigenbehaviors, are computed as the eigenvectors of the covariance matrix of behavior data, where the vectors associated to high weights represent a type of recurrent behavior, i.e., a routine. Using PCA to extract such recurrent behaviors has some advantages, compared to the simple computation of data average along each dimension. Indeed, differently from a data average, PCA allows for the extraction of multiple primary and secondary behaviors, each explaining a certain amount of data variance. Secondly, as components are evaluated along directions of maximal explained variance, they are guaranteed to accurately represent data heterogeneity, while averages can be strongly biased by the presence of recurring extreme values in the data. Finally, the amount of variance explained by each component allows a quantitative understanding of its relevance, consequently allowing for the identification and selection of primary behaviors. Accordingly to the findings presented in the previous chapter, also in this work we found that the first eigenbehavior accounts for most of the behavioral data variance of the considered user's consumption, thus we restrict our analysis to the very first eigenbehavior of each account (see the supplementary material for more information). This practically means that cumbersome hourly smart-metered time series composed of thousands of data points can be represented with one relevant daily consumption pattern, composed of only 24 data points (one for each hour of the day) and representing the most relevant consumption routine of each user, thus reducing the amount of relevant information so that it can be easily managed and represented.

5.2.2 Consumer segmentation

After data dimensionality is reduced and the principal behaviors, i.e., eigenbehaviors, are retrieved via PCA, we classify the monitored accounts into a finite number of exclusive clusters, based on the similarities of their first eigenbehavior. This step allows understanding which daily water and electricity consumption routines emerge overall the sampled accounts community. This is relevant to assess to which extent groups of accounts in all the sampled community present similar main daily consumption routines, to quantify the contribution of each group to the water and electricity demand of all the community, as well as conservation opportunities, and ultimately understand whether we can target groups of accounts with common demand management strategies. We methodologically performed this classification by sequentially implementing the t-Distributed Stochastic Neighbour Embedding (t-SNE) (Maaten and Hinton, 2008; Van Der Maaten, 2014) to emphasize data separation, and a tradi-

tional K-means clustering method (MacQueen et al., 1967). t-SNE allows for dimensionality reduction of high-dimensional objects into two-dimensional points. For our case, we applied it to the first eigenbehavior extracted to further simplify dimensionality (from 24 data points down to two data points) and enhance data separability into clusters. We then applied the K-means clustering analysis to compute the final clusters of similar water-electricity usage behaviors.

5.2.3 Factor mapping

In the last phase of our procedure we infer the most likely determinants of the first water and electricity eigenbehavior of targeted accounts clusters, among a list of candidate potential variables. Identifying water and energy demand drivers is key to customize DSM strategies and pursue water and energy efficiency by acting on such drivers. The input list of candidate drivers we considered consists of consumers' psychographic variables related to occupants' demography, households characteristics and water/energy-related attitudes and stated preferences. We extracted the most likely demand drivers out of the candidate list by mapping the latter for targeted accounts' clusters via Patient Rule Induction Method (PRIM) (Friedman and Fisher, 1999). PRIM is a factor mapping technique that has been used recently for scenario discovery under uncertainty (Kwakkel and Jaxa-Rozen, 2016) as it allows finding ranges and combinations of uncertain input variables that result in similar ranges of output variables (Kwakkel and Pruyt, 2015), thus suggesting that the selected inputs are likely to be correlated or cause the targeted output range. Given the large heterogeneity of psychographic variables overall the surveyed accounts, we applied PRIM to identify which psychographic variables, and in which range, are likely to describe consumers belonging to targeted accounts clusters (e.g., high consumers) among the ones identified after consumer segmentation (see Section 5.2.2). As we show with our results in the next section, based on the quantitative outcomes of the whole methodology it is then possible to propose recommendations for the design of customized and coordinated DSM strategies aimed at water-electricity conservation or peak shifting.

5.3 Case study

In this study, we perform user profiling and customer segmentation on a dataset consisting, after cleaning and processing, of hourly water and electricity consumption data gathered for 1107 residential accounts in South California during the period June, 28th - December, 8th, 2015. Each account is therefore associated to 4,584 hourly readings (191 days), with the full dataset including 5,074,488 data points. We only consider total water and energy household con-

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

sumption data, as no groundtruth for end-use algorithm calibration and appliance signature extraction are available at the time of this study. In addition, each account is characterized with a number of psychographic variables collected through a survey. The available variables are related to occupants' demography, households characteristics, and water/energy consumption attitudes and stated preferences. We considered 50 of these psychographic features (see the supplementary material for more information).

5.4 Numerical results

5.4.1 Analysis of coupled daily water-electricity consumption profiles

The goal of this section is analyzing the coupled daily water-electricity consumption profiles, obtained for the metered accounts through eigenbehavior analysis and then clustered to perform consumers segmentation according to the methodological steps described in Sections 5.2.1 and 5.2.2.

The first eigenbehavior of each account, representing its main water-electricity daily routine, is visualized as an individual matrix column in Figure 5.3. Each column represents the eigenbehavior coefficients, subdivided into three labelled levels of water consumption (*Zero*, *Low-medium*, and *Medium-high*) and three labelled levels of electricity consumption (*Low*, *Medium*, and *High*). Each consumption level is composed of 24 coefficients, one for each day-hour. Color intensity is proportional to coefficient magnitude: the darker the color, the higher the weight given to a certain level of consumption for the given day-hour. We labelled smart metered water data prior to eigenbehavior extraction using the following logic: *Zero demand* includes values of water consumption equal to 0 cubic feet per hour (cfh), which are over 60% of all the considered accounts hourly water readings; *Low-medium demand* includes values in the range (0, 1.67] cfh, being 1.67 cfh the amount of water used for a 5-minute shower with a flow of 2.5 gallons per minute (DeOreo et al., 2016); *Medium-high demand* includes consumption values higher than 1.67 cfh, typical, for instance, of long showers, some clothes washer programs, and outdoor uses. Being hourly electricity consumption data always positive, we labelled them as follows: *Low demand* includes values of electricity consumption lower or equal to 0.36 kWh, which is the 25-th percentile of all the accounts hourly electricity readings and is thought to contain values of base load due to plugs, stand-by appliances, or lights; *Medium demand* includes values in the range (0.36, 1.5] kWh, being 1.5 kWh the 75-th percentile of all the accounts hourly electricity readings. This range includes electricity uses like fridge and microwave; *High demand* includes consumption values higher than 1.5 kWh, e.g., air conditioning/oven. As mentioned in the previous section, we found that the first eigenbehavior explains




Figure 5.2: Histogram representing the distribution of variance explained [%] by the first eigenbehavior of each user.

most of the data variance. A histogram representing the distribution of variance explained by the first eigenbehavior is reported in Figure 5.2, showing that on average the first eigenbehavior explains more than 60% of the data variance. This metric offer a measure of how "regular" we should consider the main routine identified through eigenbehavior analysis. We found the second eigenbehavior to explain only up to 8% of data variance, meaning that secondary users' routines not explained by the first eigenbehavior are irregular and noisy, i.e., they are not really "routines".




Figure 5.3: Water-and-energy eigenbehavior coefficients and clusters. Each column of the matrix represents the coefficients of the first eigenbehavior for one user (each user is identified by means of a user ID, see x-axis). Eigenbehavior coefficients are reported for diverse labelled levels of water and energy consumption over 24 hours (see y-axis). Color intensity is proportional to coefficient magnitude, as represented in the colorbar on the right side. Eigenbehaviors are clustered in 18 groups based on their similarity: clusters are identified by the labels P1, P2, ..., P18 on top of the figure and the vertical dashed lines.

Figure 5.3 shows that there is a variety of eigenbehavior within the considered consumers' community. In particular, we identified 18 clusters of profiles through k-means clustering and silhouette assessment. These profiles are labelled with P1, P2, ..., P18 in the figure, each grouping accounts with similar eigenbehaviors. Among the profiles we can identify typical behaviors which differ in terms of weight intensity. For instance, P1 shows high weights to medium-high water and electricity consumption, meaning P1 includes high consumers, i.e., probably big houses with high number of occupants or water and/or electricity intensive end-uses like swimming pools, gardens and air conditioning systems. On the other extreme P17 includes very low consumers, i.e., single-occupant small houses/vacation houses. Profiles as P14 stand in the middle, with the highest weights during day hours given to medium levels of water and electricity consumption. Also, profiles that differ with respect to their shape during day hours can be identified: while most of the profiles show peaks in water use during early morning and evening hours, and higher electricity consumption during afternoon hours, profiles like P4 and P18 show overall low water consumption and high electricity consumption concentrated in working hours, probably because of some home business. Overall, results show that the majority of consumers do not use water during night time (dark weights to *Zero water consumption* between midnight and 6 am) and large part of day time, confirming the overall 60% of zero hourly water consumption we found in the raw data. Indeed, we can see from the empirical distributions of hourly water consumption represented in Figure 5.4 for each profile P_i that many profile show high percentage of monitored hours without positive water consumption. This can be easily explained by the fact that generally there is no water consumption when people sleep or not at home (e.g., during working hours).

Conversely, about 50% of the profiles (P1 to P9) show high weights to medium and high electricity consumption allover day and night time. The empirical distributions of hourly electricity consumption represented in Figure 5.5 for each profile P_i confirm that hourly electricity consumption is generally positive for over 95% of the time, meaning there is almost always a positive base load of energy consumption. This high base load in working and night hours might be due air conditioning systems, programmable appliances (such as clothes dryers) or many appliances always plugged in with *on* or *stand-by* mode. Despite this difference in eigenbehavior weights, in general profiles mostly weighting medium-high electricity consumption (e.g., P1 to P5) also show the highest weights for medium and high water consumption, and viceversa for low-consuming profiles like P15 or P17. Yet, the key finding is that there is not a clear matching between water and energy eigenbehaviors in terms of time-of-day usage and consumption peaks, thus we can conclude there are not evident cause-effect mechanisms between residential water and energy end-uses in our case study.

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs




Figure 5.4: Empirical Cumulative Distribution Function of hourly water consumption data [cfh], for each profile identified in the clustering phase.

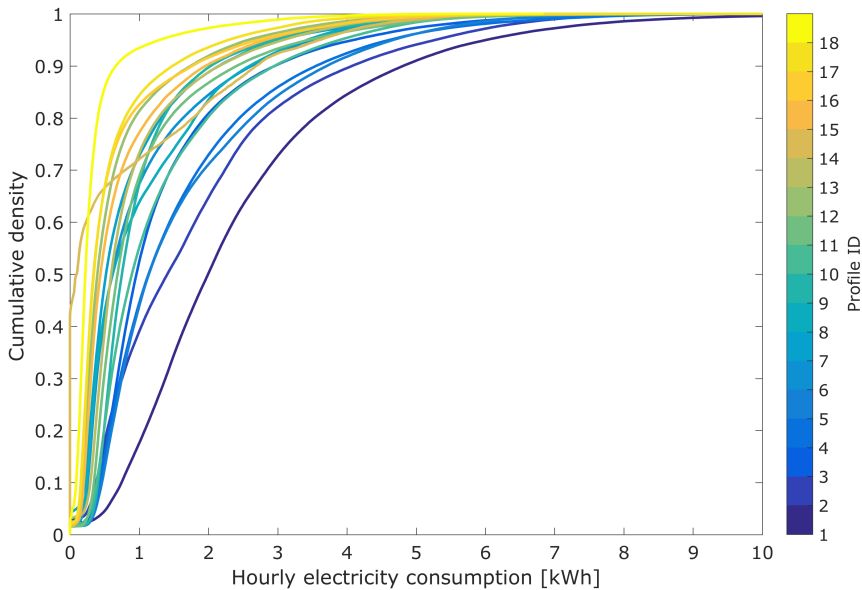


Figure 5.5: Empirical Cumulative Distribution Function of hourly electricity consumption data [kWh], for each profile identified in the clustering phase.

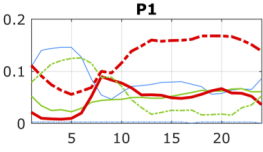
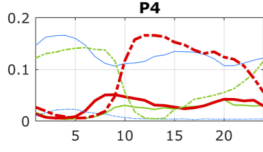
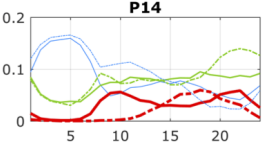
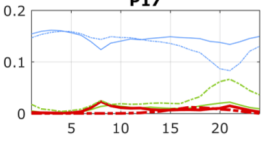
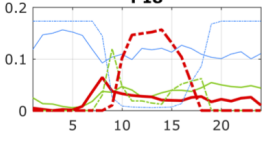
Indeed, this is enforced by the fact that water is heated by gas in most of the considered accounts, as survey data reveal. Consequently, water consumption is not the direct driver of electricity consumption, and viceversa, thus we can say that both demands are likely to be driven by other features describing houses and occupants, together with behavioural factors and attitudes.

We provide a better visualization of the different water-electricity eigenbehaviors in Figure 5.6, where the median eigenbehavior weights of each profile P1, P2, ..., P18 are represented with solid (dashed) colored lines for the 3 water (electricity) consumption levels defined above, and profiles are sorted in decreasing order of total water consumption during the observation period. Median profiles confirm the overall mismatch between water and electricity time-of-use and consumption peaks, i.e., the solid red line indicating high water consumption shows two peaks, the highest in early morning hours (6-10 am) and a second in the evening (6-11 pm), while the dashed red line, indicating high electricity consumption shows a bigger peak that usually starts in late morning/early afternoon hours and then lasts till late evening hours. Also, the main heterogeneous typologies of behaviors identified in Figure 5.3 can be recognized, and we better characterize them in Table 5.1 also with respect to DSM interventions. Finally, behavioral differences emerge among high-consuming water profiles, denoting two kinds of water-intensive behaviors. Profiles such as P1 and P3 describe users that regularly use high amounts of water during day hours, demonstrated by the larger values of the solid red (high usage) and green (medium usage) lines relative to the solid blue line (low usage) during peak hours. In contrast, profiles like P2 and P4 describe the behavior of users that only occasionally use high amounts of water during day hours (evidenced by the larger values of the solid blue line relative to the solid green and red lines for all hours), but total water consumption during the monitored period comparable to that of P1 and P3. These characterizations of profiles in terms of time-of-use of water and electricity and frequency of high demands are essential to support the design of DSM strategies based on pricing mechanisms. Indeed, increasing block tariffs, i.e., rates with increase in the unit price of water as more water is consumed (Spang and Loge, 2015), penalizing high water and electricity consumption may be effective at incentivizing conservation by consumers characterized by profiles of "flat" high consumption like P4, while dynamic hourly pricing schemes may be a better choice for reducing the "peakiness" of consumption (and thus reducing utilities costs during peak hours (Beal et al., 2016)) evidenced by such profiles as P3. In practice, offering only selected pricing schemes to different groups of customers can be unfeasible, as that might cause inequalities and, in general, pricing involves social and political challenges. Still, utilities can exploit the information of different customer profiles to design a common pricing scheme (e.g., either just block tariffs or

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

hourly dynamic pricing) suitable to target the majority of critical profiles and maximize the overall benefit in terms of conservation and cost savings. Moreover, they can adapt the selected pricing scheme over time, based on observed behavioral changes and modifications of consumption intensity and timing.

Table 5.1: Characterization of main heterogeneous typologies of behaviors.

Profile ID	Profile Name	Characteristics and implications for DM
	High-regular water-electricity consumer	<ul style="list-style-type: none"> Users that regularly use high amounts of water and electricity during day hours, with high "peakiness" in peak hours (red solid and dashed lines are higher than blue and green lines during peak hours). Dynamic hourly water and electricity pricing schemes may reduce the "peakiness" of their demand. Increasing block tariffs, efficient devices, and conservation feedbacks can incentivize overall conservation.
	High-occasional water consumer	<ul style="list-style-type: none"> Users that only occasionally use high amount of water during day hours (green and red solid lines are dominated by the blue solid line), with a total water demand close to that of <i>High-regular consumers</i> and less marked peaks than <i>High-regular consumers</i>. High level of electricity consumption. Increasing block water tariffs, efficient devices, and conservation feedbacks can incentivize overall water conservation.
	Average water-electricity consumer	<ul style="list-style-type: none"> Mostly, medium levels of water (solid green line) and energy consumption (dashed green line) during daytime hours, with seldom higher peaks during peak hours. We do not consider them main priority targets for DSM interventions. Still, customized feedbacks and efficient devices can help reducing water-electricity waste.
	Low consumer	<ul style="list-style-type: none"> Low consuming profiles, with high coefficients to <i>Zero</i> water consumption (blue solid line) and <i>Low</i> electricity consumption (blue dashed line). We do not consider them a priority targets for DSM interventions.
	Daytime consumer	<ul style="list-style-type: none"> Low water consumption and high electricity consumption concentrated in working hours. Perhaps, users with this profile run a home business. Given the overall low water and electricity consumption, we do not consider them main priority targets for DSM interventions.

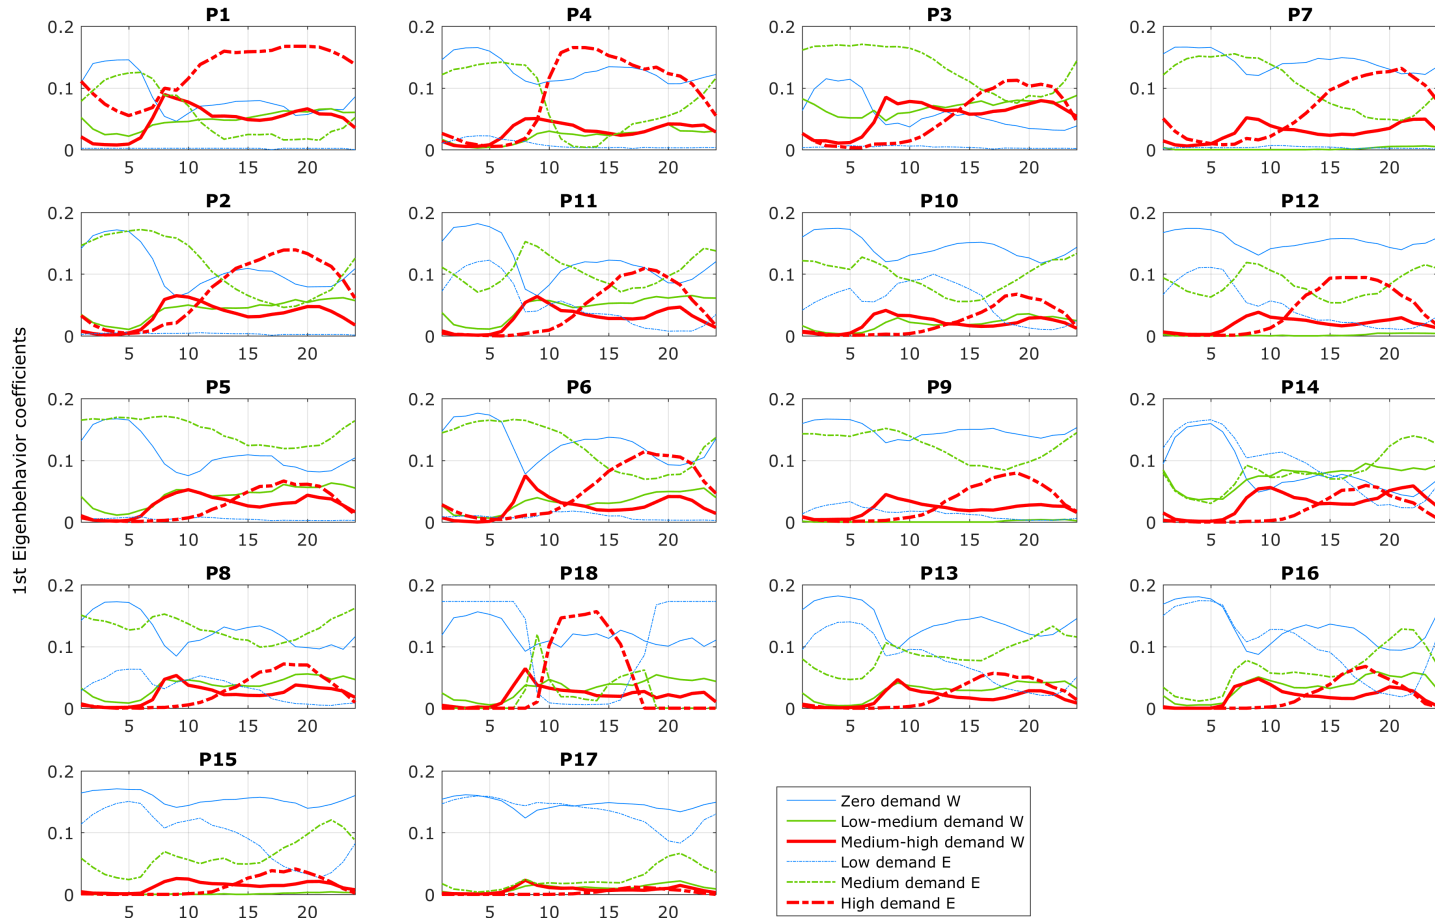


Figure 5.6: Median weights of the first eigenbehavior for each profile P1, P2, ..., P18. Profiles are sorted in decreasing order of total per-capita water consumption during the observation period (from top-left to bottom-right). Solid lines refer to water, dashed lines to electricity. Colors refer to different consumption labels.

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

5.4.2 Recommendations for the design of customized demand-side management strategies

In order to formulate further recommendations to DSM, in this section we cross-analyze consumption profiles and users' psychographic variables through factor mapping analysis. Given that most of the candidate variables for factor mapping were retrieved as consumers' psychographic variables through a water-related survey to consumers, results better characterize water demand. Still, they also provide important hints in terms of energy demand management.

5.4.2.1 Water-Energy conservation programs

We observed that average water demand and average electricity demand of each profile show a linear correlation equal to 0.93 (see Figure 5.7), despite the main daily water consumption routines have been shown to differ in terms of time-of-use from daily electricity consumption routines. Based on this finding, we identified a set of profiles characterized by large water and electricity consumption. Such profiles, visualized with a dark and light blue boxes in Figure 5.7, can be considered high-priority targets for water and energy conservation actions, as they contribute to the large share of water-electricity demand.

Targeting the very top consuming profile (P1), up to the top 5 profiles (P1, P4, P7, P3, P2), we run factor mapping analysis according to the methodology of Section 5.2, in order to explore the potential drivers of their demands. Results are summarized in Table 5.2, under the *Water conservation* demand management (DM) scenario.

Results show that the restricted dimension, i.e., the main factor identified as critical for explaining the profile with higher water and energy consumption (P1, *Top 1* in the table) is the presence of either a swimming pool, or a hot tub, or both, as highlighted in bold in Table 5.2 under column *Box limits*. Around 75% of the users belonging to profile P1 declared to own a swimming pool, a hot tub, or both, while the remaining declared to own "Neither" a pool or a hot tub. This finding is in line with previous studies (Mayer and DeOreo, 1999) that demonstrated that not only swimming pools are one of the most consuming end-uses - homes with swimming pools use more than twice as much water outdoors than homes without swimming pools - but also that those households with a swimming pool on average showed a more consuming attitude for other end-uses. In our case, the selected restricted dimension can also explain part of the high electricity demand of targeted profiles, as pools and hot tubs require electricity for pumping and filtration. Since a specific outdoor end-use is selected as potential driver of demand, results suggests that demand management strategies such as feedback customized on that specific end use, as well as block tariffs for high consumption or season and time-dependent restrictions regulating the use of




Figure 5.7: Average per-household daily water and electricity consumption for each consumer cluster. Each scatter point represents a cluster of consumers. Clusters falling inside the dark and light blue boxes are potential targets for water and energy conservation interventions.

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

water for outdoor end-uses during scarcity periods can be designed to pursue water (and partially energy) conservation. Extending the analysis up to the *Top 5* profiles, the previous result is confirmed and pool/hot tub presence is still selected. Additionally, demographic aspects (i.e., number of house occupants) emerge as relevant, as we also hypothesised in Section 5.4.1. Water demand increases with the number of house occupants (Cominola et al., 2015b), meaning that some top-consuming accounts show high demand just because of their size and the overlapping routines of many users. In that case, users can be targeted with customized feedbacks aimed at increasing efficiency in water and energy use and exploiting economies of scale in the household (e.g., running the clothes washer only when full).

We repeated the factor mapping analysis only on those users who do not have a pool or a hot tub to see whether factor mapping selects other variables as demand drivers. Interestingly, attitudes and subjective perceptions become relevant, because users belonging to the top consuming clusters declared a medium-to-low sensitivity towards water price and a medium-to-low environmental interest towards saving water. While keeping into account that there is often a difference between what consumers declare (stated preference) and how they actually behave (revealed preferences) (Beal et al., 2013), this information is essential to DSM, as programs based on water-energy awareness and education towards education can be effective to improve the efficiency in these users' water-energy consumption behaviors.

5.4.2.2 Water peak shifting programs

The objective of DSM is not only restricted to pursue reductions in water and energy demand, but also reducing utilities costs during peak hours (Beal et al., 2016), especially in periods with high resources availability, is a goal of DSM. Thinking of a scenario in which water utilities are interested in reducing costs due to morning peak water use, we repeated factor mapping to target those profiles contributing to most of water demand between 6 and 10 am.

While results mainly confirm the findings discussed in Section 5.4.2.1, two interesting results emerge when considering the *Top 4* profiles targeted for peak shifting (see the last line in Table 5.2). First, the influence of swimming pool is enforced as the critical range of that variable is now restricted to having “Both a pool and a hot tub” and “Pool only”, because the peak flow of a swimming pool, or its combination with the peak flow of a hot tub is larger than that of only a hot tub. Secondly, relevance of personal attitudes is further enhanced, as targeted users declared a medium-to-low trust in believing their water utility helps them understanding their water use. This is an important indication for education-based DSM actions and for the development of services aimed at providing consumers' with accessible and detailed information about their

Table 5.2: Factor mapping results for customers' clusters to target for water demand management programs.

DM scenario	Target clusters	Restricted dimensions	Box limits
Water conservation	Top 1	Pool/hot tub presence	Both a pool and a hot tub Pool only Hot tub only Neither
Water conservation	Top 5	Pool/hot tub presence Number of house occupants	Both a pool and a hot tub Pool only Hot tub only Neither 3-8
Water conservation	Top 1 w/o pool	Interest in spending less on water bill	Strongly disagree Somewhat disagree Neither agree or disagree Somewhat agree Strongly agree
Water conservation	Top 3 w/o pool	Saving water helps the environment	Strongly disagree Somewhat disagree Neither agree or disagree Somewhat agree Strongly agree
Peak shifting	Top 4	Pool/hot tub presence Number of house occupants Believe water utility makes it easy understanding their water use	Both a pool and a hot tub Pool only Hot tub only Neither 3-8 Strongly disagree Somewhat disagree Neither agree or disagree Somewhat agree Strongly agree

consumption, as well as incentives to promote efficient behaviors (Novak et al., 2016).

Overall, our findings prove that the consumer segmentation procedure we contribute is suitable to explore differences between typical residential water and electricity demand patterns, as well as providing useful insights to coordinated water-energy demand management programs based on resources conservation or demand peak shifting. On the customer side, results demonstrate that both actual household features and personal attitudes or perceptions contribute to characterize water and energy consumers' behaviors.

5.5 Discussion

In this chapter we integrated the users' behavioral models we developed and proposed along this thesis, in order to exploit the information about users' characterization they provide to design customized water and energy demand-side management strategies on targeted groups of users. The integrated procedure we proposed combines three methodological phases, i.e., (i) eigenbehavior ex-

5. Customer segmentation analysis to infer residential water-electricity demand patterns and drivers and inform customized demand-side management programs

traction for the identification of typical consumers profiles from smart metered hourly water-energy time series, (ii) clustering for customer segmentation based on profile similarities, and (iii) factor mapping to infer the potential determinants of targeted profiles to inform the design of customized DSM. We applied our methodology to a case study including over 1000 residential accounts in South California, monitored at hourly resolution for over six months, with the goals of exploring heterogeneity of typical residential water-electricity demand profiles, interpreting them in terms of users' behavior, performing customer segmentation and providing insights to coordinated water-energy DSM for targeted groups of high-consumers.

Results demonstrated that our procedure is suitable to capture the heterogeneity in water-electricity consumption profiles and routines over a community of users, highlighting differences between daily time-of-use of water and electricity, and allowing for the characterization of accounts based on psychographic and behavioral factors. Moreover, factor mapping has been shown to be able to discriminate among the factors most likely to influence water and energy demand for targeted groups of users.

In terms of key implications for DSM on our specific case study, firstly we found that daily water and energy consumption profiles are not much correlated in terms of daily demand pattern shape. This is mainly because water is mostly heated by gas in South California, thus there is not a clear causal nexus between water and energy demand, and viceversa. This means that demand management strategies aimed at changing the daily demand patterns (e.g., hourly dynamic tariffs) should be differentiated between the two sources. At the same time, despite their different load shape, we found a correlation between the total amount of water and energy consumers' use. This is suggesting that common groups of users can be targeted for water-energy DSM.

Finally, we found that both objective factors (e.g., presence of a swimming pool), as well as subjective and personal ones (e.g., attitudes toward resources conservation) are relevant in order to distinguish users based on their consumption habits and routines. The type of factors emerging from the factor mapping phase is essential in order to customize demand management strategies in order to really act on each user's consumption drivers. Yet, we should also keep into account that there is often a difference between declared and observed users' preferences, attitudes and, in general, behaviors (Beal et al., 2013). This can affect the effectiveness of demand management strategies, therefore repeated data analysis and behavioral change monitoring are needed, in order to adapt such strategies depending on the achievements reached over time. Overall, we think that the data mining procedure we propose in this research represents a promising tool advancing traditional methods for inferring residential water and electricity consumer behaviors. Indeed, the growing development of

AMI systems worldwide is opening up new opportunities for demand modeling and management in terms of data sampling resolution and data frequency of acquisition. Yet, the burden of dataset dimensionality and the challenges of information extraction should be tackled with data mining techniques. In this regard, we demonstrated the suitability of our methodology to address such challenges. Further applications of our methodology on bigger datasets should be performed, in order to test its scalability to city- or utility-scale analysis.

Further research should also include the application of our analytic framework to communities of users monitored for longer than 6 months, in order to include in our analysis the effect of meteorological conditions and seasonality on consumers' behavior. Having a larger number of entire weekly period monitored at high resolution would also allow distinguishing differences in water and energy consumption behavior between weekdays and weekends. Further, the availability of consumption data measured or estimated at the end-use level would allow a more detailed characterization and differentiation of users' behaviors, as well as the assessment of the influence of users' psychographic features on specific end-uses. Finally, a joint behavioral experimental study comprising both a data analysis phase and a practical phase with the actual implementation of demand-size management strategies on controlled groups of accounts, designed on the basis of the output provided by our methodology, would be an essential contribution to prove and measure the actual suitability of our framework to support DSM.

6

Conclusions and future research

Research on urban water demand modelling and management in the residential sector is undergoing a rapid evolution, accounting for changing conditions in both demand and resource availability, technological opportunities offered by smart meters, big data, and synergies with other research, operation, and business sectors, like energy.

Following these premises, the main goal of this thesis was advancing data analysis and mathematical models to extract information on water consumers' behavior out of smart-metered data, ultimately informing and proposing recommendations to customized demand management.

As an outcome of our research, we developed novel methodologies, data mining, and modelling tools contributing to water and energy end-uses characterization, water users profiling and behavioral modeling, and targeting of water and energy consumers for the design of personalized WDMS and recommendations. More in detail, the main deliverables of this research are the following. First, a review of more than 130 studies, published between 1990 and 2015, on high and low resolution residential water demand modelling and management. Apart from being the first comprehensive review on this topic, it proposes also a methodological framework for the classification and comparison of studies within that field, and a "roadmap" listing the main research challenges that application of smart metering technologies will face in the near future. Second, two novel Non-Intrusive Load Monitoring algorithms demonstrated to achieve high electric power and water disaggregation performance. The application of such algorithms on both real-world and synthetically generated water and energy data provided insights on the robustness of the two methods to signal noise, data resolution, as well as their portability between

6. Conclusions and future research

the two application fields. Third, two novel modelling tools to infer water consumers' habits and routines, as well as identify the most relevant determinants of their water consuming or saving behaviors, at the household level. Tests of both tools on high and low resolution water consumption data from two real-world case studies in Switzerland and Australia demonstrated their suitability for inferring typical profiles of water consumption, accurately predicting water consumption profiles as a function of consumers' socio-psychographic information, and performing customers segmentation. Finally, a three-phase data-mining procedure composed of data dimensionality reduction, customer segmentation, and factor mapping. Application of this integrated procedure on water and electricity data from over 1000 accounts in South California demonstrated the suitability of our procedure for identifying heterogeneous groups of users to be potentially targeted for the design of customized water and energy demand management strategies either for conservation or peak-shifting objectives.

Overall, we think that the methods proposed and developed in this research, as well as its findings, constitute important progress for an effective and efficient exploitation of smart metering data to develop models of water-energy users' behavior at the household scale, and advance the customization of demand-side management strategies.

Beside the specific findings of each phase of this research, which we discussed in detail at the end of each chapter, we can list some general take-home conclusions and recommendations for future research and technological development. We found that the development of accurate descriptive and predictive models of water and energy consumers' behaviors benefits from consumption data metered with high spatial and temporal granularity, as well as from the availability of data characterizing users' socio-demographic characteristics, habits, and attitudes. Coupling proper data mining techniques to such data availability allows extracting concise information to support utilities and management decision processes. In general, data gathering campaigns can be costly and time-consuming, and only a few experiences around the world allowed gathering consistent dataset for model development, calibration, and testing. Therefore, developing automatic and transparent procedures which integrate data from many sources can cost-effectively increase data availability, increase the frequency of updates of such data, and increase users' acceptability. From a technological point of view, our research suggests that obtaining very high accuracy on end-use disaggregation would require data sampling resolutions in the order of few seconds or few minutes. Due to the limitations posed by battery life-cycle, data transmission and storage capabilities, and the computational needs of data analysis and modelling methodologies, slightly lower resolutions can still provide useful end-use information (e.g., to spot potential areas

for water and energy conservation, or opportunities for appliance retrofitting). Utilities, meter producers, and data scientists should keep into account these technological issues when planning new smart metering deployments, and balance data requirements with the physical limitations posed by smart metering technology. This is also relevant in relation to those aspects related to quasi real-time network management and control, which have not been addressed in this thesis, but constitute part of the core business of utilities. Finally, we believe that inter-sectoral synergies, benchmarking, and collaborations, especially between the water and energy field, is essential to create standardized and portable procedures and algorithms, explore interconnections and links between such sectors, and promote studies supporting the activities of multi-utilities and coordinated management.

Starting from the outcomes of this thesis, follow-up research should focus on the following aspects:

- the water and energy end-use algorithms we proposed should be further tested onto real-world high resolution water data, as well as tested in terms of scalability on a large dataset comprising several hundreds, or thousands users;
- coupled applications of our behavioural models with real-world implementations of water and energy demand management strategies should be performed in collaboration with utilities, in order to assess the suitability of such models to detect changes in users' behavior in time and monitor the effectiveness and progress of the implemented management actions;
- coordinated test of our models on consistent data from water-energy users from different geographical, climate, and social areas can improve our understanding of context influence on water-energy users' attitudes and habits, as well as the comparison of suitable context-specific demand management strategies;
- the integration of the models and data mining tools we developed into a unique product can open up new opportunities for large-scale testing, final deployment and adoption by utilities and water-energy management authorities.

Bibliography

- Abdallah, A., Rosenberg, D., 2014. Heterogeneous residential water and energy linkages and implications for conservation and management. *Journal of Water Resources Planning and Management* 140 (3), 288–297.
- Abed-Meraim, K., Qiu, W., Hua, Y., 1997. Blind system identification. *Proceedings of the IEEE* 85 (8), 1310–1322.
- Abrahamse, W., Steg, L., Vlek, C., Rothengatter, T., 2007. The effect of tailored information, goal setting, and tailored feedback on household energy use, energy-related behaviors, and behavioral antecedents. *Journal of Environmental Psychology* 27 (4), 265–276.
- Agudelo-Vera, C., Blokker, E., Büscher, C., Vreeburg, J., 2014. Analysing the dynamics of transitions in residential water consumption in the netherlands. *Water Science & Technology: Water Supply* 14 (5), 717–727.
- Allcott, H., Greenstone, M., 2012. Is there an energy efficiency gap? Tech. rep., National Bureau of Economic Research.
- Altunkaynak, A., Ozger, M., Cakmakci, M., 2005. Water consumption prediction of istanbul city by using fuzzy logic approach. *Water Resources Management* 19 (5), 641–654.
- Alvisi, S., Franchini, M., Marinelli, A., 2007. A short-term, pattern-based model for water-demand forecasting. *Journal of Hydroinformatics* 9 (1), 39.
- Aly, A. H., Wanakule, N., 2004. Short-term forecasting for urban water consumption. *Journal of water resources planning and management* 130 (5), 405–410.
- Amenta, V., Tina, G. M., 2015. Load Demand Disaggregation Based on Simple Load Signature and User's Feedback. *Energy Procedia* 83, 380–388.
- Anda, M., Le Gay Brereton, F., Brennan, J., E., P., 2013. Smart Metering Infrastructure for Residential Water Efficiency: Results of a Trial in a Behavioural Change Program in Perth, Western Australia.
- Aquacraft Inc., 2011. Albuquerque single-family water use efficiency and retrofit study. Rep. Prepared for the Albuquerque Bernalillo County Water.
- Arbués, F., Garcia-Valiñas, M. Á., Martinez-Espiñeira, R., 2003. Estimation of residential water demand: a state-of-the-art review. *The Journal of Socio-Economics* 32 (1), 81–102.
- Arregui, F., Cabrera, E., Cobacho, R., 2006. Integrated water meter management. IWA Publishing London.
- Assembly, U. G., 2015. Transforming our world: the 2030 agenda for sustainable development. New York: United Nations.

Bibliography

- Balling, R. C., Gober, P., 2007. Climate Variability and Residential Water Use in the City of Phoenix, Arizona. *Journal of Applied Meteorology and Climatology* 46 (7), 1130–1137.
- Balling, R. C., Gober, P., Jones, N., 2008. Sensitivity of residential water consumption to variations in climate: an intraurban analysis of Phoenix, Arizona. *Water Resources Research* 44 (10).
- Baranski, M., Voss, J., 2004. Genetic algorithm for pattern detection in NILM systems. In: *Systems, Man and Cybernetics, 2004 IEEE International Conference on*. Vol. 4. IEEE, pp. 3462–3468.
- Barker, S., Kalra, S., Irwin, D., Shenoy, P., 2014. NILM redux: The case for emphasizing applications over accuracy. In: *NILM-2014 Workshop*.
- Batra, N., Kelly, J., Parson, O., Dutta, H., Knottenbelt, W., Rogers, A., Singh, A., Srivastava, M., 2014. NILMTK: An open source toolkit for non-intrusive load monitoring. In: *Proceedings of the 5th international conference on Future energy systems*. pp. 265–276.
- Beal, C., Stewart, R., 2014. Identifying Residential Water End Uses Underpinning Peak Day and Peak Hour Demand. *Journal of Water Resources Planning and Management* 140 (7).
- Beal, C., Stewart, R. A., Huang, T., Rey, E., 2011a. South east queensland residential end use study. *Journal of the Australian Water Association* 38 (1), 80–84.
- Beal, C., Stewart, R. A., Spinks, A., Fielding, K., 2011b. Using smart meters to identify social and technological impacts on residential water consumption. *Water Science and Technology-Water Supply* 11 (5), 527–533.
- Beal, C. D., Gurung, T. R., Stewart, R. A., 2016. Demand-side management for supply-side efficiency: Modeling tailored strategies for reducing peak residential water demand. *Sustainable Production and Consumption* 6, 1–11.
- Beal, C. D., Makki, A., Stewart, R. A., 2014. What does rebounding water use look like? An examination of post-drought and post-flood water end-use demand in Queensland, Australia. *Water Science and Technology-Water Supply* 14 (4), 561–568.
- Beal, C. D., Stewart, R. A., Fielding, K., 2013. A novel mixed method smart metering approach to reconciling differences between perceived and actual residential end use water consumption. *Journal of Cleaner Production* 60, 116–128.
- Bennett, C., Stewart, R. A., Beal, C. D., 2013. Ann-based residential water end-use demand forecasting model. *Expert Systems with Applications* 40 (4), 1014–1023.
- Benzi, F., Anglani, N., Bassi, E., Frosini, L., 2011. Electricity Smart Meters Interfacing the Households. *IEEE Transactions on Industrial Electronics* 58 (10), 4487–4494.
- Berk, R., Cooley, T., LaCivita, C., Parker, S., Sredl, K., Brewer, M., 1980. Reducing consumption in periods of acute scarcity: the case of water. *Soc. Sci. Res.:(United States)* 9 (2).
- Bernard, T., Marx, M., 2016. Unsupervised Learning Algorithm using multiple Electrical Low and High Frequency Features for the task of Load Disaggregation. In: *Proceedings of the 3rd International Workshop on NILM*.
- Blokker, E., Vreeburg, J., van Dijk, J., 2010. Simulating residential water demand with a stochastic end-use model. *Journal of Water Resources Planning and Management* 136 (1), 19–26.
- Bogost, I., 2007. *Persuasive games: The expressive power of videogames*. Mit Press.
- Bonabeau, E., 2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America* 99 (Suppl 3), 7280–7287.

- Bonfigli, R., Squartini, S., Fagiani, M., Piazza, F., 2015. Unsupervised algorithms for non-intrusive load monitoring: An up-to-date overview. In: Environment and Electrical Engineering (EEEIC), 2015 IEEE 15th International Conference on. IEEE, pp. 1175–1180.
- Borisova, T., Useche, P., 2013. Exploring the effects of extension workshops on household water-use behavior. *HortTechnology* 23 (5), 668–676.
- Bousquet, F., Le Page, C., 2004. Multi-agent simulations and ecosystem management: a review. *Ecological Modelling* 176 (3-4), 313–332.
- Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., Stewart, R., 2013. Intelligent metering for urban water: A review. *Water* 5 (3), 1052–1081.
- Brennan, D., Tapsuwan, S., Ingram, G., 2007. The welfare costs of urban outdoor water restrictions. *Australian Journal of Agricultural and Resource Economics* 51 (3), 243–261.
- Britton, T., Cole, G., Stewart, R., Wiskar, D., 2008. Remote diagnosis of leakage in residential households. *Journal of Australian Water Association* 35 (6), 89–93.
- Britton, T. C., Stewart, R. A., O'Halloran, K. R., 2013. Smart metering: enabler for rapid and effective post meter leakage identification and water loss management. *Journal of Cleaner Production* 54 (0), 166–176.
- Brooks, D. B., 2006. An operational definition of water demand management. *International Journal of Water Resources Development* 22 (4), 521–528.
- Brookshire, D. S., Burness, H. S., Chermak, J. M., Krause, K., 2002. Western urban water demand. *Nat. Resources J.* 42, 873.
- Butner, R. S., Reid, D. J., Hoffman, M., Sullivan, G. P., Blanchard, J., 2013. Non-Intrusive Load Monitoring Assessment: Literature Review and Laboratory Protocol. Pacific Northwest National Laboratory.
- Cahill, R., Lund, J., 2012. Residential water conservation in Australia and California. *Journal of Water Resources Planning and Management* 139 (1), 117–121.
- Cahill, R., Lund, J., DeOreo, W. B., Medellín-Azuara, J., 2013. Household water use and conservation models using Monte Carlo techniques. *Hydrology and Earth System Sciences Discussions* 10 (4), 4869–4900.
- Cardell-Oliver, R., 2013a. Discovering water use activities for smart metering. In: Intelligent Sensors, Sensor Networks and Information Processing, 2013 IEEE Eighth International Conference on. IEEE, pp. 171–176.
- Cardell-Oliver, R., 2013b. Water use signature patterns for analyzing household consumption using medium resolution meter data. *Water Resources Research* 49 (12), 8589–8599.
- Cardell-Oliver, R., Peach, G., 2013. Making sense of smart metering data. *Australian Water Association Water Journal* 4 (40), 124–128.
- Cardell-Oliver, R., Wang, J., Gigney, H., 2016. Smart meter analytics to pinpoint opportunities for reducing household water use. *Journal of Water Resources Planning and Management* 142 (6), 04016007.
- Carragher, B. J., Stewart, R. A., Beal, C. D., 2012. Quantifying the influence of residential water appliance efficiency on average day diurnal demand patterns at an end use level: A precursor to optimised water service infrastructure planning. *Resources Conservation and Recycling* 62, 81–90.
- Carrie Armel, K., Gupta, A., Shrimali, G., Albert, A., 2013. Is disaggregation the holy grail of energy efficiency? the case of electricity. *Energy Policy* 52, 213–234.

Bibliography

- Cawley, G. C., Talbot, N. L., Girolami, M., 2007. Sparse multinomial logistic regression via bayesian l1 regularisation. *Advances in neural information processing systems* 19, 209.
- Chang, H., Parandvash, G. H., Shandas, V., 2010. Spatial Variations of Single-Family Residential Water Consumption in Portland, Oregon. *Urban Geography* 31 (7), 953–972.
URL <http://dx.doi.org/10.2747/0272-3638.31.7.953>
- Chen, D., Irwin, D., Shenoy, P., Albrecht, J., 2014. Combined Heat and Privacy: Preventing Occupancy Detection from Smart Meters. In: *2014 Ieee International Conference on Pervasive Computing and Communications (percom)*, pp. 208–215.
- Chou, J.-S., Yutami, I. G. A. N., 2014. Smart meter adoption and deployment strategy for residential buildings in Indonesia. *Applied Energy* 128, 336 – 349.
- Chu, J., Wang, C., Chen, J., Wang, H., 2009. Agent-based residential water use behavior simulation and policy implications: A case-study in beijing city. *Water resources management* 23 (15), 3267–3295.
- Clifford, E., Coakley, D., Curry, E., Degeler, V., Costa, A., Messervey, T., Van Andel, S.-J., Van de Giesen, N., Kouroupetroglou, C., Mink, J., et al., 2014. Interactive water services: The waternomics approach. *Procedia Engineering* 89, 1058–1065.
- Colak, I., Fulli, G., Sagirolglu, S., Yesilbudak, M., Covrig, C.-F., 2015. Smart grid projects in Europe: Current status, maturity and future scenarios . *Applied Energy* 152, 58 – 70.
- Cole, G., O'Halloran, K., Stewart, R. A., 2012. Time of use tariffs: implications for water efficiency. *Water Science and Technology-Water Supply* 12 (1), 90–100.
- Cole, G., Stewart, R. A., 2013. Smart meter enabled disaggregation of urban peak water demand: precursor to effective urban water planning. *Urban Water Journal* 10 (3), 174–194.
- Collins, R., Kristensen, P., Thyssen, N., 2009. Water resources across Europe-confronting water scarcity and drought. Office for Official Publications of the European Communities.
- Cominola, A., Giuliani, M., Castelletti, A., Abdallah, A., Rosenberg, D., 11 – 14 July 2016a. Developing a stochastic simulation model for the generation of residential water end-use demand time series. In: *Proceedings of the 8th International Congress on Environmental Modeling and Software (iEMSs 2016)*. Toulouse (France).
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2015a. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environmental Modelling & Software* 72, 198–214.
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., 2017. A hybrid signature-based iterative disaggregation algorithm for non-intrusive load monitoring. *Applied Energy* 185, 331–344.
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A., Anda, M., 28 June – 3 July 2015b. Modelling residential water consumers' behaviors by feature selection and feature weighting. In: *Proceedings of the 36th IAHR World Congress*. The Hague (the Netherlands).
- Cominola, A., Moro, L., Riva, A., Giuliani, M., Castelletti, A., 11 – 14 July 2016b. Profiling residential water users' routines by eigenbehavior modelling. In: *Proceedings of the 8th International Congress on Environmental Modeling and Software (iEMSs 2016)*. Toulouse (France).
- Cominola, A., Spang, E., Giuliani, M., Castelletti, A., Loge, F., Lund, J., 2016c. Residential water and electricity customer segmentation analysis to inform demand-side management programs. *Environmental Science & Technology* (in preparation).
- Commission, N. W., et al., 2013. National performance report 2011–12: urban water utilities.

- Corbella, H. M., Pujol, D. S., 2009. What lies behind domestic water use? a review essay on the drivers of domestic water consumption. *Boletín de la Asociación de Geógrafos Españoles* 50, 297–314.
- Cordell, D., Robinson, J., Loh, M., 2003. Collecting residential end use data from primary sources: do's and don't's. *Proceedings of Efficient 2003*.
- Cosgrove, C. E., Cosgrove, W. J., 2012. The Dynamics of Global Water Futures: Driving Forces 2011-2050. Vol. 2 of *The United Nations World Water Development Report*. UNESCO.
- Costanza, E., Ramchurn, S. D., Jennings, N. R., 2012. Understanding domestic energy consumption through interactive visualisation: a field study. In: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, pp. 216–225.
- Cover, T. M., Thomas, J. A., 2012. *Elements of information theory*. John Wiley & Sons.
- Curry, E., Degeler, V., Clifford, E., Coakley, D., Costa, A., van Andel, S., van de Giesen, N., Kouroupetroglou, C., Messervey, T., Mink, J., et al., 2014. Linked water data for water information management. In: *11th International Conference on Hydroinformatics (HIC)*, New York, New York, USA.
- Dalhuisen, J. M., Florax, R. J., de Groot, H. L., Nijkamp, P., 2003. Price and income elasticities of residential water demand: a meta-analysis. *Land Economics* 79 (2), 292–308.
- Dandy, G., Nguyen, T., Davies, C., 1997. Estimating residential water demand in the presence of free allowances. *Land Economics*, 125–139.
- Darby, S., 2006. The effectiveness of feedback on energy consumption: A review for DEFRA of the literature on metering, billing and direct displays. *Environmental Change Institute*, Oxford, UK.
- Davies, K., Doolan, C., van den Honert, R., Shi, R., 2014. Water-saving impacts of smart meter technology: An empirical 5 year, whole-of-community study in sydney, australia. *Water Resources Research* 50 (9), 7348–7358.
- DeOreo, W., Mayer, P., Dziegielwski, B., Kiefer, J., 2016. *Residential Uses of Water 2016*. Water Research Foundation, Denver, CO.
- DeOreo, W. B., 2011. Analysis of water use in new single family homes. By Aquacraft. For Salt Lake City Corporation and US EPA.
- DeOreo, W. B., Heaney, J. P., Mayer, P. W., 1996. Flow trace analysis to assess water use. *AWWA* 88 (1), 79–90.
- DeOreo, W. B., Mayer, P. W., 1994. Project Report: A Process Approach for Measuring Residential Water Use and Assessing Conservation. Tech. rep.
- DeOreo, W. B., Mayer, P. W., 2000. The end uses of hot water in single family homes from flow trace analysis. *Seattle Public Utilities and US EPA*, Seattle.
- DeOreo, W. B., Mayer, P. W., Martien, L., Hayden, M., Funk, A., Kramer-Duffield, M., Davis, R., Henderson, J., Raucher, B., Gleick, P., et al., 2011. California single-family water use efficiency study. Rep. Prepared for the California Dept. of Water Resources.
- Desley, V., Laurie, B., Peter, M., 2013. The effectiveness of energy feedback for conservation and peak demand: a literature review. *Open Journal of Energy Efficiency* 2013.
- Desmedt, J., Vekemans, G., Maes, D., 2009. Ensuring effectiveness of information to influence household behaviour. *Journal of Cleaner Production* 17 (4), 455–462.

Bibliography

- Dong, M., Meira, P. C., Xu, W., Chung, C., 2013. Non-intrusive signature extraction for major residential loads. *Smart Grid, IEEE Transactions on* 4 (3), 1421–1430.
- Donoho, D., 2006. Compressed sensing. *IEEE Trans. on Information Theory* 52 (4), 1289–1306.
- Duda, R. O., Hart, P. E., et al., 1973. *Pattern classification and scene analysis*. Vol. 3. Wiley New York.
- Duke, J. M., Ehemann, R. W., Mackenzie, J., 2002. The distributional effects of water quantity management strategies: A spatial analysis. *The Review of Regional Studies* 32 (1), 19–35.
- Eagle, N., Pentland, A. S., 2009. Eigenbehaviors: Identifying structure in routine. *Behavioral Ecology and Sociobiology* 63 (7), 1057–1066.
- EEA, 2001. Sustainable water use in Europe – Part 2: Demand Management. Tech. Rep. Environmental issue report No.19, EEA - European Environment Agency.
- Ehrhardt-Martinez, K., Donnelly, K. A., Laitner, S., et al., 2010. Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities. American Council for an Energy-Efficient Economy Washington, DC.
- Elafoudi, G., Stankovic, L., Stankovic, V., 2014. Power disaggregation of domestic smart meter readings using dynamic time warping. In: *Communications, Control and Signal Processing (ISCCSP), 2014 6th International Symposium on*. IEEE, pp. 36–39.
- Elhamifar, E., Sastry, S., 2015. Energy Disaggregation via Learning Powerlets and Sparse Coding. In: *AAAI*. pp. 629–635.
- Escrivà Bou, À., 2016. The water-energy nexus: a bottom-up approach for basin-wide management. Ph.D. thesis.
- Escriva-Bou, A., Lund, J. R., Pulido-Velazquez, M., 2015a. Modeling residential water and related energy, carbon footprint and costs in california. *Environmental Science & Policy* 50, 270–281.
- Escriva-Bou, A., Lund, J. R., Pulido-Velazquez, M., 2015b. Optimal residential water conservation strategies considering related energy in california. *Water Resources Research*.
- Espey, M., Espey, J., Shaw, W. D., 1997. Price elasticity of residential demand for water: A meta-analysis. *Water Resources Research* 33 (6), 1369–1374.
- Espinoza, M., Joye, C., Belmans, R., De Moor, B., 2005. Short-term load forecasting, profile identification, and customer segmentation: a methodology based on periodic time series. *IEEE Transactions on Power Systems* 20 (3), 1622–1630.
- Evans, R. P., Blotter, J. D., Stephens, A. G., 2004. Flow rate measurements using flow-induced pipe vibration. *Journal of fluids engineering* 126 (2), 280–285.
- Farinaccio, L., Zmeureanu, R., 1999. Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses. *Energy and Buildings* 30 (3), 245 – 259.
- Fernández, C., Manyà, F., Mateu, C., Sole-Mauri, F., 2014. Modeling energy consumption in automated vacuum waste collection systems. *Environmental Modelling & Software* 56 (0), 63 – 73, thematic issue on Modelling and evaluating the sustainability of smart solutions.
- Fielding, K. S., Spinks, A., Russell, S., McCrea, R., Stewart, R. A., Gardner, J., 2013. An experimental test of voluntary strategies to promote urban water demand management. *Journal of environmental management* 114, 343–351.

- Figueiredo, M., Ribeiro, B., de Almeida, A., Feb 2014. Electrical signal source separation via nonnegative tensor factorization using on site measurements in a smart home. *Instrumentation and Measurement, IEEE Transactions on* 63 (2), 364–373.
- Fischer, C., 2008. Feedback on household electricity consumption: a tool for saving energy? *Energy efficiency* 1 (1), 79–104.
- Fox, C., McIntosh, B., Jeffrey, P., 2009. Classifying households for water demand forecasting using physical property characteristics. *Land Use Policy* 26 (3), 558 – 568.
- Friedman, J. H., Fisher, N. I., 1999. Bump hunting in high-dimensional data. *Statistics and Computing* 9 (2), 123–143.
- Froehlich, J., Findlater, L., Ostergren, M., Ramanathan, S., Peterson, J., Wragg, I., Larson, E., Fu, F., Bai, M., Patel, S., Landay, J. A., 2012. The design and evaluation of prototype eco-feedback displays for fixture-level water usage data. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. ACM, New York, NY, USA, pp. 2367–2376.
- Froehlich, J., Larson, E., Saba, E., Campbell, T., Atlas, L., Fogarty, J., Patel, S., 2011. A longitudinal study of pressure sensing to infer real-world water usage events in the home. In: *Pervasive Computing*. Springer, pp. 50–69.
- Froehlich, J. E., Larson, E., Campbell, T., Haggerty, C., Fogarty, J., Patel, S. N., 2009. Hydrosense: Infrastructure-mediated single-point sensing of whole-home water activity. In: *Proceedings of the 11th International Conference on Ubiquitous Computing*. Ubicomp '09. ACM, New York, NY, USA, pp. 235–244.
- Froes Lima, C. A., Portillo Navas, J. R., 2012. Smart metering and systems to support a conscious use of water and electricity. *Energy* 45 (1), 528–540.
- Fuchs, J., 2004. On sparse representations in arbitrary redundant bases. *IEEE Trans. on Information Theory* 50 (6), 1341–1344.
- Fullerton, T. M., Elias, A., 2004. Short-term water consumption dynamics in El Paso, Texas. *Water Resources Research* 40 (8).
- Gabaldón, A., Ortiz-García, M., Molina, R., Valero-Verdú, S., 2014. Disaggregation of the electric loads of small customers through the application of the hilbert transform. *Energy Efficiency*, 1–18.
- Gaiardelli, E., 2015. Development of a flexible synthetic generator of residential water consumption traces at the end-use level. Master's thesis, Politecnico di Milano, Italy.
- Gaiser, K., Stroeve, P., 2014. The impact of scheduling appliances and rate structure on bill savings for net-zero energy communities: Application to West Village. *Applied Energy* 113, 1586 – 1595.
- Galán, J., López-Paredes, A., Del Olmo, R., 2009. An agent-based model for domestic water management in valladolid metropolitan area. *Water Resources Research* 45 (5).
- Galelli, S., Castelletti, A., 2013. Tree-based iterative input variable selection for hydrological modeling. *Water Resources Research* 49 (7), 4295–4310.
- Galelli, S., Humphrey, G. B., Maier, H. R., Castelletti, A., Dandy, G. C., Gibbs, M. S., 2014. An evaluation framework for input variable selection algorithms for environmental data-driven models. *Environmental Modelling & Software* 62, 33–51.
- Garcia, S., Thomas, A., 2001. The structure of municipal water supply costs: Application to a panel of french local communities. *Journal of Productivity Analysis* 16 (1), 5–29.

Bibliography

- Gato, S., 2006. Forecasting urban residential water demand. Ph.D. thesis, School of Civil, Environmental and Chemical Engineering, RMIT University.
- Gato, S., Jayasuriya, N., Roberts, P., 2007. Temperature and rainfall thresholds for base use urban water demand modelling. *Journal of Hydrology* 337 (3), 364 – 376.
- Gato-Trinidad, S., Jayasuriya, N., Roberts, P., 2011. Understanding urban residential end uses of water. *Water Science & Technology* 64 (1).
- Geller, E. S., 2002. The challenge of increasing proenvironment behavior. In: *Handbook of environmental psychology*. Wiley: New York, NY, USA, Ch. 34, pp. 525–540.
- Geller, H., Harrington, P., Rosenfeld, A. H., Tanishima, S., Unander, F., 2006. Polices for increasing energy efficiency: Thirty years of experience in OECD countries. *Energy Policy* 34 (5), 556 – 573, hong Kong Editorial Board meeting presentationsHong Kong Editorial Board meeting presentations.
- George, E. I., 2000. The variable selection problem. *Journal of the American Statistical Association* 95 (452), 1304–1308.
- Gerland, P., Raftery, A. E., evÅ ikova, H., Li, N., Gu, D., Spoorenberg, T., Alkema, L., Fosdick, B. K., Chunn, J., Lalic, N., Bay, G., Buettner, T., Heilig, G. K., Wilmoth, J., 2014. World population stabilization unlikely this century. *Science* 346 (6206), 234–237.
- Ghahramani, Z., Jordan, M. I., 1997. Factorial hidden Markov models. *Machine learning* 29 (2-3), 245–273.
- Giacomini, M., Berglund, E., 2015. Complex adaptive modeling framework for evaluating adaptive demand management for urban water resources sustainability. *Journal of Water Resources Planning and Management*, 04015024.
- Gilg, A., Barr, S., 2006. Behavioural attitudes towards water saving? evidence from a study of environmental actions. *Ecological Economics* 57 (3), 400–414.
- Giri, S., Bergés, M., 2015. An energy estimation framework for event-based methods in Non-Intrusive Load Monitoring. *Energy Conversion and Management* 90, 488–498.
- Giuliani, M., Mossina, J., 2016. Reducing the intrusiveness of energy and water end-use disaggregation via social media and users interactions. In: *Proceedings of the 1st International Workshop on the Social Web for Environmental and Ecological Monitoring, SWEEM 2016*.
- Giurco, D. P., White, S. B., Stewart, R. A., 2010. Smart metering and water end-use data: conservation benefits and privacy risks. *Water* 2 (3), 461–467.
- Gleick, P. H., Haasz, D., Henges-Jeck, C., Srinivas, V., Wolff, G., Cushing, K. K., Mann, A., 2003. Waste not, want not: The potential for urban water conservation in California. Pacific Institute for Studies in Development, Environment, and Security Oakland, CA.
- Goncalves, H., Ocneanu, A., Berges, M., Fan, R., 2011. Unsupervised disaggregation of appliances using aggregated consumption data. In: *The 1st KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*.
- Grafton, R. Q., Ward, M. B., 2008. Prices versus Rationing: Marshallian Surplus and Mandatory Water Restrictions. *Economic Record* 84, S57–S65.
- Grafton, R. Q., Ward, M. B., To, H., Kompas, T., 2011. Determinants of residential water consumption: Evidence and analysis from a 10-country household survey. *Water Resources Research* 47 (8).

- Griffin, R. C., Chang, C., 1991. Seasonality in community water demand. *Western Journal of Agricultural Economics* 16 (2), 207–217.
- Gullo, F., Ponti, G., Tagarelli, A., Ruffolo, M., Labate, D., et al., 2009. Low-voltage electricity customer profiling based on load data clustering. In: *Proceedings of the 2009 International Database Engineering & Applications Symposium*. ACM, pp. 330–333.
- Gurung, T. R., Stewart, R. A., Beal, C. D., Sharma, A. K., 2015. Smart meter enabled water end-use demand data: enhanced infrastructure planning of contemporary supply networks platform for the urban water supply networks. *Journal of Cleaner Production* 87, 642–654.
- Gurung, T. R., Stewart, R. A., Sharma, A. K., Beal, C. D., 2014. Smart meters for enhanced water supply network modelling and infrastructure planning. *Resources Conservation and Recycling* 90, 34–50.
- Guyon, I., Elisseeff, A., 2003. An introduction to variable and feature selection. *Journal of machine learning research* 3 (Mar), 1157–1182.
- Guyon, I., Weston, J., Barnhill, S., Vapnik, V., 2002. Gene selection for cancer classification using support vector machines. *Machine learning* 46 (1-3), 389–422.
- Harou, J., Garrone, P., Rizzoli, A., Maziotis, A., Castelletti, A., Fraternali, P., Novak, J., Wissmann-Alves, R., Ceschi, P., 2014. Smart metering, water pricing and social media to stimulate residential water efficiency: Opportunities for the smarth2o project. *Procedia Engineering* 89, 1037–1043.
- Hart, G. W., 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE* 80 (12), 1870–1891.
- Heinrich, M., 2007. Water End Use and Efficiency Project (WEEP) – Final Report. BRANZ Study Report 159.
- Hensher, D., Shore, N., Train, K., 2006. Water supply security and willingness to pay to avoid drought restrictions. *Economic Record* 82 (256), 56–66.
- Hilty, L. M., Aebischer, B., Rizzoli, A. E., 2014. Modeling and evaluating the sustainability of smart solutions. *Environmental Modelling & Software* 56 (0), 1 – 5, thematic issue on Modelling and evaluating the sustainability of smart solutions.
- Hilty, L. M., Arnfalk, P., Erdmann, L., Goodman, J., Lehmann, M., Wäger, P. A., 2006. The relevance of information and communication technologies for environmental sustainability—a prospective simulation study. *Environmental Modelling & Software* 21 (11), 1618 – 1629, *environmental Informatics*.
- Homwongs, C., Sastri, T., Foster III, J. W., 1994. Adaptive forecasting of hourly municipal water consumption. *Journal of water resources planning and management* 120 (6), 888–905.
- House-Peters, L. A., Chang, H., 2011. Urban water demand modeling: Review of concepts, methods, and organizing principles. *Water Resources Research* 47 (5).
- Howarth, D., Butler, S., 2004. Communicating water conservation: how can the public be engaged? *Water Science & Technology: Water Supply* 4 (3), 33–44.
- Howe, C. W., 1982. The impact of price on residential water demand: Some new insights. *Water Resources Research* 18 (4), 713–716.
- Howe, C. W., Linaweaver, F. P., 1967. The impact of price on residential water demand and its relation to system design and price structure. *Water Resources Research* 3 (1), 13–32.
- IEA, O., 2012. *World energy outlook 2012*. Vienna: IEA.

Bibliography

- Inman, D., Jeffrey, P., 2006. A review of residential water conservation tool performance and influences on implementation effectiveness. *Urban Water Journal* 3 (3), 127–143.
- Janmaat, J., 2013. Spatial patterns and policy implications for residential water use: An example using Kelowna, British Columbia. *Water Resources and Economics* 1, 3–19.
- Jolliffe, I., 2002. *Principal Component Analysis*. Springer, New York, N.Y.
- Jolliffe, I., 2002. *Principal component analysis*. Wiley Online Library.
- Jorgensen, B., Graymore, M., O’Toole, K., 2009. Household water use behavior: An integrated model. *Journal of Environmental Management* 91 (1), 227 – 236.
- Kanakoudis, V., 2002. Urban water use conservation measures. *Aqua- Journal of Water Supply: Research and Technology* 51 (3), 153–163.
- Kanta, L., Zechman, E., 2014. Complex Adaptive Systems Framework to Assess Supply-Side and Demand-Side Management for Urban Water Resources. *Journal of Water Resources Planning and Management* 140 (1), 75–85.
- Kelly, J., Knottenbelt, W., 2015. 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*. ACM, pp. 55–64.
- Kenney, D. S., Goemans, C., Klein, R., Lowrey, J., Reidy, K., 2008. Residential Water Demand Management: Lessons from Aurora, Colorado. *JAWRA Journal of the American Water Resources Association* 44 (1), 192–207.
- Kenney, D. S., Klein, R. A., Clark, M. P., 2004. Use and Effectiveness of Municipal Water Restrictions During Drought in Colorado. *JAWRA Journal of the American Water Resources Association* 40 (1), 77–87.
- Kenny, J. F., Barber, N. L., Hutson, S. S., Linsey, K. S., Lovelace, J. K., Maupin, M. A., 2009. Estimated use of water in the United States in 2005. US Geological Survey Reston, VA.
- Kim, H., Marwah, M., Arlitt, M. F., Lyon, G., Han, J., 2011. Unsupervised Disaggregation of Low Frequency Power Measurements. In: *SDM*. Vol. 11. SIAM, pp. 747–758.
- Kim, Y., Schmid, T., Charbiwala, Z. M., Friedman, J., Srivastava, M. B., 2008. Nawms: nonintrusive autonomous water monitoring system. In: *Proceedings of the 6th ACM conference on Embedded network sensor systems*. ACM, pp. 309–322.
- Klein, G., Krebs, M., Hall, V., O’Brien, T., Blevins, B., 2005. California’s water-energy relationship. California Energy Commission.
- Kobus, C. B., Klaassen, E. A., Mugge, R., Schoormans, J. P., 2015. A real-life assessment on the effect of smart appliances for shifting households’ electricity demand. *Applied Energy* 147, 335 – 343.
- Kofinas, D., Mellios, N., Papageorgiou, E., Laspidou, C., 2014. Urban water demand forecasting for the island of Skiathos. *Procedia Engineering* 89, 1023–1030.
- Kolter, J. Z., Batra, S., Ng, A. Y., 2010. Energy disaggregation via discriminative sparse coding. In: *NIPS*. pp. 1153–1161.
- Kolter, J. Z., Jaakkola, T., 2012. Approximate inference in additive factorial HMMs with application to energy disaggregation. In: *International conference on artificial intelligence and statistics*. pp. 1472–1482.

- Kolter, J. Z., Johnson, M. J., 2011. Redd: A public data set for energy disaggregation research. In: Workshop on Data Mining Applications in Sustainability (SIGKDD), San Diego, CA.
- Kossieris, P., Panayiotakis, A., Tzouka, K., Gerakopoulou, P., Rozos, E., Makropoulos, C., 2014. An elearning approach for improving household water efficiency. *Procedia Engineering* 89, 1113–1119.
- Kowalski, M., Marshallsay, D., 2003. A system for improved assessment of domestic water use components. International Water Association.
- Kowalski, M., Marshallsay, D., 2005. Using measured microcomponent data to model the impact of water conservation strategies on the diurnal consumption profile. *Water Science & Technology: Water Supply* 5.
- Kramers, A., Höjer, M., Lövehagen, N., Wangel, J., 2014. Smart sustainable cities—exploring ict solutions for reduced energy use in cities. *Environmental Modelling & Software* 56 (0), 52 – 62, thematic issue on Modelling and evaluating the sustainability of smart solutions.
- Kwac, J., Flora, J., Rajagopal, R., 2014. Household energy consumption segmentation using hourly data. *IEEE Transactions on Smart Grid* 5 (1), 420–430.
- Kwakkel, J. H., Jaxa-Rozen, M., 2016. Improving scenario discovery for handling heterogeneous uncertainties and multinomial classified outcomes. *Environmental Modelling & Software* 79, 311–321.
- Kwakkel, J. H., Pruyt, E., 2015. Using system dynamics for grand challenges: The esdma approach. *Systems Research and Behavioral Science* 32 (3), 358–375.
- Laniak, G. F., Olchin, G., Goodall, J., Voinov, A., Hill, M., Glynn, P., Whelan, G., Geller, G., Quinn, N., Blind, M., Peckham, S., Reaney, S., Gaber, N., Kennedy, R., Hughes, A., 2013. Integrated environmental modeling: A vision and roadmap for the future. *Environmental Modelling & Software* 39 (0), 3 – 23, thematic Issue on the Future of Integrated Modeling Science and Technology.
- Lapidou, C., 2014. Ict and stakeholder participation for improved urban water management in the cities of the future.
- Lee, M., Tansel, B., Balbin, M., 2011. Influence of residential water use efficiency measures on household water demand: A four year longitudinal study. *Resources, Conservation and Recycling* 56 (1), 1 – 6.
- Lee, S.-J., Wentz, E. A., 2008. Applying Bayesian Maximum Entropy to extrapolating local-scale water consumption in Maricopa County, Arizona. *Water Resources Research* 44 (1).
- Lee, S.-J., Wentz, E. A., 2010. Effects of urban spatial structure, sociodemographics, and climate on residential water consumption in Hillsboro, Oregon. *Journal of American Water Resources Association* 46 (3), 461–472.
- Letcher, R. A. K., Jakeman, A. J., Barreteau, O., Borsuk, M. E., ElSawah, S., Hamilton, S. H., Henriksen, H. J., Kuikka, S., Maier, H. R., Rizzoli, A. E., van Delden, H., Voinov, A. A., 2013. Selecting among five common modelling approaches for integrated environmental assessment and management. *Environmental Modelling & Software* 47 (0), 159 – 181.
- Liang, J., Ng, S. K., Kendall, G., Cheng, J. W., 2010. Load signature study - Part II: Disaggregation framework, simulation, and applications. *IEEE Transactions on Power Delivery* 25 (2), 561–569.
- Liao, J., Elafoudi, G., Stankovic, L., Stankovic, V., 2014. Non-intrusive appliance load monitoring using low-resolution smart meter data. In: *Smart Grid Communications (SmartGridComm)*, 2014 IEEE International Conference on. IEEE, pp. 535–540.

Bibliography

- Ligtenberg, A., van Lammeren, R. J. A., Bregt, A. K., Beulens, A. J. M., 2010. Validation of an agent-based model for spatial planning: A role-playing approach. *Computers, Environment and Urban Systems* 34 (5), 424–434.
- Likas, A., Vlassis, N., Verbeek, J., 2003. The global k-means clustering algorithm. *Pattern recognition* 36 (2), 451–461.
- Little, L., Flynn, J., 2012. *Big Data, Smart Meters & Australian Water Utilities*. Tech. rep.
- Liu, B., Luan, W., Yu, Y., 2016. A Fully Unsupervised Appliance Modelling Framework for NILM. In: *Proceedings of the 3rd International Workshop on NILM*.
- Liu, H., Setiono, R., 1995. Chi2: Feature selection and discretization of numeric attributes. In: *2012 IEEE 24th International Conference on Tools with Artificial Intelligence*. IEEE Computer Society, pp. 388–388.
- Loh, M., Coghlan, P., Australia, W., 2003. Domestic water use study: In Perth, Western Australia, 1998–2001. Water Corporation.
- Loureiro, D., Alegre, H., Coelho, S., Martins, A., Mamade, A., 2014. A new approach to improve water loss control using smart metering data. *Water Science & Technology: Water Supply* 14 (4), 618–625.
- Lyman, R. A., 1992. Peak and off-peak residential water demand. *Water Resources Research* 28 (9), 2159–2167.
- Maaten, L. v. d., Hinton, G., 2008. Visualizing data using t-sne. *Journal of Machine Learning Research* 9 (Nov), 2579–2605.
- MacQueen, J., et al., 1967. Some methods for classification and analysis of multivariate observations. In: *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*. Vol. 1. Oakland, CA, USA., pp. 281–297.
- Maggioni, E., 2015. Water demand management in times of drought: What matters for water conservation. *Water Resources Research*.
- Magiera, E., Froelich, W., 2014. Integrated support system for efficient water usage and resources management (iss-ewatus). *Procedia Engineering* 89, 1066–1072.
- Maidment, D. R., Parzen, E., 1984. Cascade model of monthly municipal water use. *Water Resources Research* 20 (1), 15–23.
- Maier, H. R., Dandy, G. C., 2000. Neural networks for the prediction and forecasting of water resources variables: a review of modelling issues and applications. *Environmental modelling & software* 15 (1), 101–124.
- Makki, A. A., Stewart, R. A., Beal, C. D., Panuwatwanich, K., 2015. Novel bottom-up urban water demand forecasting model: Revealing the determinants, drivers and predictors of residential indoor end-use consumption. *Resources, Conservation and Recycling* 95, 15–37.
- Makki, A. A., Stewart, R. A., Panuwatwanich, K., Beal, C., 2013. Revealing the determinants of shower water end use consumption: enabling better targeted urban water conservation strategies. *Journal of Cleaner Production* 60, 129–146.
- Makonin, S., Popowich, F., 2015. Nonintrusive load monitoring (NILM) performance evaluation. *Energy Efficiency* 8 (4), 809–814.
- Makonin, S., Popowich, F., Bartram, L., Gill, B., Bajic, I., 2013. AMPds: a public dataset for load disaggregation and eco-feedback research. In: *Electrical Power and Energy Conference (EPEC)*. pp. 1–6.

- Matos, C., Teixeira, C. A., Bento, R., Varajao, J., Bentes, I., Jan. 2014. An exploratory study on the influence of socio-demographic characteristics on water end uses inside buildings. *Science of The Total Environment* 466-467, 467–474.
- Mauch, L., Barsim, K. S., Yang, B., 2016. How well can HMM model load signals. In: *Proceedings of the 3rd International Workshop on NILM*.
- Mayer, P. W., DeOreo, W. B., 1995. Process approach for measuring residential water use and assessing conservation effectiveness. AWWA.
- Mayer, P. W., DeOreo, W. B., 1999. Residential end uses of water. American Water Works Association.
- Mayer, P. W., Deoreo, W. B., Lewis, D. M., 2000. Seattle home water conservation study: The impacts of high efficiency plumbing fixture retrofits in single-family homes. Tech. rep., Seattle Public Utilities and USEPA, Boulder, Colorado.
- Mayer, P. W., Deoreo, W. B., Towler, E., Lewis, D. M., 2003. Residential indoor water conservation study: Evaluation of high efficiency indoor plumbing fixture retrofits in single-family homes in the east bay municipal utility district (edmund) service area. Tech. rep., The United States Environmental Protection Agency.
- Mayer, P. W., DeOreo, W. B., Towler, E., Martien, L., Lewis, D., 2004. Tampa water department residential water conservation study: the impacts of high efficiency plumbing fixture retrofits in single-family homes. A Report Prepared for Tampa Water Department and the United States Environmental Protection Agency.
- McDonald, R., Douglas, I., Grimm, N., Hale, R., Revenga, C., Gronwall, J., Fekete, B., 2011a. Implications of fast urban growth for freshwater provision. *Ambio* 40 (5), 437.
- McDonald, R., Green, P., Balk, D., Fekete, B., Revenga, C., Todd, M., Montgomery, M., 2011b. Urban growth, climate change, and freshwater availability. *Proceedings of the National Academy of Sciences* 108 (15), 6312–6317.
- McDonald, R. I., Weber, K., Padowski, J., Flörke, M., Schneider, C., Green, P. A., Gleeson, T., Eckman, S., Lehner, B., Balk, D., Boucher, T., Grill, G., Montgomery, M., 2014. Water on an urban planet: Urbanization and the reach of urban water infrastructure. *Global Environmental Change* 27 (0), 96 – 105.
- McIntyre, T., 2008. Data retention in Ireland: Privacy, policy and proportionality. *Computer Law & Security Review* 24 (4), 326–334.
- Mead, N., Aravinthan, V., 2009. Investigation of household water consumption using smart metering system. *Desalination and Water Treatment* 11 (1-3), 115–123.
- Miaou, S.-P., 1990. A class of time series urban water demand models with nonlinear climatic effects. *Water Resources Research* 26 (2), 169–178.
- Molino, B., Rasulo, G., Tagliatalata, L., 1996. Forecast model of water consumption for naples. *Water resources management* 10 (4), 321–332.
- Molinos-Senante, M., 2014. Water rate to manage residential water demand with seasonality: peak-load pricing and increasing block rates approach. *Water Policy* 16 (5), 930–944.
- Mori, M., Tezuka, K., Tezuka, H., Dec. 24 2004. Ultrasonic flow meter, flow measurement method, and computer program. US Patent App. 10/584,318.
- Moro, A., Riva, L., 2016. Inferring and clustering residential water consumers' routines by eigenbehavior modeling. Master's thesis, Politecnico di Milano, Italy.

Bibliography

- Morsali, H., Shekarabi, S., Ardekani, K., Khayami, H., Fereidunian, A., Ghassemian, M., Lesani, H., May 2012. Smart plugs for building energy management systems. In: Smart Grids (ICSG), 2012 2nd Iranian Conference on. pp. 1–5.
- Mueller, J. A., Kimball, J. W., 2016. An Accurate Method of Energy Use Prediction for Systems with Known Composition. In: Proceedings of the 3rd International Workshop on NILM.
- Mühlhäuser, M., Gurevych, I., Global, I., 2008. Handbook of research on ubiquitous computing technology for real time enterprises. Information Science Reference.
- Müller, M., 2007. Dynamic time warping. Information retrieval for music and motion, 69–84.
- Nasseri, M., Moeini, A., Tabesh, M., 2011. Forecasting monthly urban water demand using extended kalman filter and genetic programming. Expert Systems with Applications 38 (6), 7387–7395.
- Neenan, B., Hemphill, R. C., 2008. Societal benefits of smart metering investments. The Electricity Journal 21 (8), 32 – 45.
- Newborough, M., Probert, S., 1990. Intelligent automatic electrical-load management for networks of major domestic appliances. Applied Energy 37 (2), 151 – 168.
- Nguyen, K., Stewart, R., Zhang, H., 2013b. An intelligent pattern recognition model to automate the categorisation of residential water end-use events. Environmental Modelling & Software 47, 108–127.
- Nguyen, K. A., Stewart, R. A., Zhang, H., 2014. An autonomous and intelligent expert system for residential water end-use classification. Expert Systems with Applications 41 (2), 342–356.
- Nguyen, K. A., Zhang, H., Stewart, R. A., 2013a. Development of an intelligent model to categorise residential water end use events. Journal of Hydro-environment Research 7 (3), 182 – 201.
- Niese, A., TrÄschel, M., Sonnenschein, M., 2014. Designing dependable and sustainable smart grids—how to apply algorithm engineering to distributed control in power systems. Environmental Modelling & Software 56 (0), 37 – 51, thematic issue on Modelling and evaluating the sustainability of smart solutions.
- Nizar, A., Dong, Z., 2009. Identification and detection of electricity customer behaviour irregularities. In: Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES. IEEE, pp. 1–10.
- Novak, J., Melenhorst, M., Micheel, I., Pasini, C., Fraternali, P., Rizzoli, A. E., 10 – 14 July 2016. Behaviour change and incentive modelling for water saving: first results from the SmartH2O project. In: Proceedings of the 8th International Congress on Environmental Modelling and Software. Toulouse (France).
- Olmstead, S. M., Michael Hanemann, W., Stavins, R. N., 2007. Water demand under alternative price structures. Journal of Environmental Economics and Management 54 (2), 181–198.
- Olmstead, S. M., Stavins, R. N., 2009. Comparing price and nonprice approaches to urban water conservation. Water Resources Research 45 (4).
- Oracle, 2009. Smart Metering for Water Utilities. Tech. rep.
- Parson, O., Ghosh, S., Weal, M., Rogers, A., 2014. An unsupervised training method for non-intrusive appliance load monitoring. Artificial Intelligence 217, 1–19.
- Pereira, L., Quintal, F., Barreto, M., Nunes, N. J., 2013. Understanding the limitations of eco-feedback: A one-year long-term study. In: Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data. Springer, pp. 237–255.

- Pérez, R., Cugueró, M.-A., Cugueró, J., Sanz, G., 2014. Accuracy assessment of leak localisation method depending on available measurements. *Procedia Engineering* 70, 1304–1313.
- Perez, R., Sanz, G., Puig, V., Quevedo, J., Cuguero Escofet, M. A., Nejjari, F., Meseguer, J., Cembrano, G., Mirats Tur, J. M., Sarrate, R., 2014. Leak localization in water networks: A model-based methodology using pressure sensors applied to a real network in barcelona [applications of control]. *Control Systems, IEEE* 34 (4), 24–36.
- Peschiera, G., Taylor, J. E., Siegel, J. A., 2010. Response–relapse patterns of building occupant electricity consumption following exposure to personal, contextualized and occupant peer network utilization data. *Energy and Buildings* 42 (8), 1329–1336.
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A., 28 June – 3 July 2015. A convex optimization approach for automated water and energy end use disaggregation. In: *Proceedings of the 36th IAHR World Congress*. Vol. 28. The Hague (the Netherlands).
- Piga, D., Cominola, A., Giuliani, M., Castelletti, A., Rizzoli, A. E., 2016. Sparse optimization for automated energy end use disaggregation. *IEEE Transactions on Control Systems Technology* 24 (3), 1044–1051.
- Pöschacker, M., Egarter, D., Elmenreich, W., 2016. Proficiency of Power Values for Load Disaggregation. *IEEE Transactions on Instrumentation and Measurement* 65 (1), 46–55.
- Polebitski, A. S., Palmer, R. N., 2010. Seasonal residential water demand forecasting for census tracts. *Journal of Water Resources Planning and Management* 136 (1), 27–36.
- Ponce, M. V. C., Castañón, L. E. G., Cayuela, V. P., et al., 2014. Model-based leak detection and location in water distribution networks considering an extended-horizon analysis of pressure sensitivities. *Journal of Hydroinformatics* 16 (3), 649–670.
- Praskievicz, S., Chang, H., 2009. Identifying the Relationships Between Urban Water Consumption and Weather Variables in Seoul, Korea. *Physical Geography* 30 (4), 324–337.
- Qi, C., Chang, N.-B., 2011. System dynamics modeling for municipal water demand estimation in an urban region under uncertain economic impacts. *Journal of Environmental Management* 92 (6), 1628–1641.
- Quinlan, J. R., 2014. *C4. 5: programs for machine learning*. Elsevier.
- Renwick, M. E., Archibald, S. O., 1998. Demand side management policies for residential water use: who bears the conservation burden? *Land economics*, 343–359.
- Renwick, M. E., Green, R. D., 2000. Do residential water demand side management policies measure up? an analysis of eight california water agencies. *Journal of Environmental Economics and Management* 40 (1), 37 – 55.
- Reynders, E., 2012. System identification methods for (operational) modal analysis: review and comparison. *Archives of Computational Methods in Engineering* 19 (1), 51–124.
- Rezgui, Y., Zarli, A., Ellis, K. A., McCann, J. A., 2014. Optimized water demand management through intelligent sensing and analytics: The wisdom approach. In: *11th International Conference on Hydroinformatics (HIC) 2014*, New York City, USA.
- Rixon, A., Moglia, M., Burn, S., 2007. Exploring water conservation behaviour through participatory agent based modelling. *Topics on Systems Analysis for Integrated Water Resource Management*.
- Rizzoli, A., Castelletti, A., Cominola, A., Fraternali, P., dos Santos, A. D., Storni, B., Wissmann-Alves, R., Bertocchi, M., Novak, J., Micheel, I., Jun. 2014. The smarth2o project and the role of social computing in promoting efficient residential water use: a first analysis. In: *Proceedings of the 7th International Congress on Environmental Modelling and Software*.

Bibliography

- Roberts, P., 2005. Yarra Valley Water: 2004 residential end use measurement study. Tech. rep., Yarra Valley Water Melbourne, Melbourne, Australia.
- Romano, G., Salvati, N., Guerrini, A., 2014. Estimating the determinants of residential water demand in Italy. *Water* 6 (10), 2929–2945.
- Rosenberg, D., 2010. Residential water demand under alternative rate structures: Simulation approach. *Journal of Water Resources Planning and Management* 136 (3), 395–402.
- Rosenberg, D. E., Tarawneh, T., Abdel-Khaleq, R., Lund, J. R., 2007. Modeling integrated water user decisions in intermittent supply systems. *Water Resources Research* 43 (7), W07425.
- Rowlands, I. H., Reid, T., Parker, P., 2014. Research with disaggregated electricity end-use data in households: review and recommendations. *Wiley Interdisciplinary Reviews: Energy and Environment*. URL <http://dx.doi.org/10.1002/wene.151>
- Russell, S., Fielding, K., 2010. Water demand management research: A psychological perspective. *Water Resources Research* 46 (5), W05302.
- Ruzzelli, A. G., Nicolas, C., Schoofs, A., O'Hare, G. M., 2010. Real-time recognition and profiling of appliances through a single electricity sensor. In: *Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010 7th Annual IEEE Communications Society Conference on*. IEEE, pp. 1–9.
- Ryan, S., Long, R., Lauf, D., Ledbetter, M., Reeves, A., 2010. Energy star water heater market profile. Efficiency Sells, US Department of Energy.
- Sakoe, H., Chiba, S., 1978. Dynamic programming algorithm optimization for spoken word recognition. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 26 (1), 43–49.
- Sanderson, M. L., Yeung, H., 2002. Guidelines for the use of ultrasonic non-invasive metering techniques. *Flow Measurement and Instrumentation* 13 (4), 125–142.
- Savić, D., Vamvakieridou-Lyroudia, L., Kapelan, Z., 2014. Smart meters, smart water, smart societies: The iwidget project. *Procedia Engineering* 89, 1105–1112.
- Schneider, M., Whitlatch, E., 1991. User-specific water demand elasticities. *Journal of Water Resources Planning and Management* 117 (1), 52–73.
- Schultz, P. W., Messina, A., Tronu, G., Limas, E. F., Gupta, R., Estrada, M., 2014. Personalized Normative Feedback and the Moderating Role of Personal Norms A Field Experiment to Reduce Residential Water Consumption. *Environment and Behavior*, 0013916514553835.
- SDU, 2011. Urban water management plan 2010. Tech. rep., City of Sacramento, Sacramento, California.
- Shao, H., Marwah, M., Ramakrishnan, N., 2012. A temporal motif mining approach to unsupervised energy disaggregation. In: *Proceedings of the 1st International Workshop on Non-Intrusive Load Monitoring*, Pittsburgh, PA, USA. Vol. 7.
- Shoham, Y., Leyton-Brown, K., 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Princeton University Press.
- Singh, S., Gulati, M., Majumdar, A., 2016. Greedy Deep Disaggregating Sparse Coding. In: *Proceedings of the 3rd International Workshop on NILM*.
- SJESD, 2011. Urban water management plan. Tech. rep., City of San Jose, Sacramento, California.
- Sonderlund, A. L., Smith, J. R., Hutton, C., Kapelan, Z., 2014. Using Smart Meters for Household Water Consumption Feedback: Knowns and Unknowns. *Procedia Engineering* 89, 990–997.

- Spang, E. S., Loge, F. J., 2015. A high-resolution approach to mapping energy flows through water infrastructure systems. *Journal of Industrial Ecology* 19 (4), 656–665.
- Srinivasan, D., Ng, W., Liew, A., 2006. Neural-network-based signature recognition for harmonic source identification. *Power Delivery, IEEE Transactions on* 21 (1), 398–405.
- Steg, L., Vlek, C., 2009. Encouraging pro-environmental behaviour: An integrative review and research agenda. *Journal of Environmental Psychology* 29 (3), 309–317.
- Stewart, R. A., Willis, R., Giurco, D., Panuwatwanich, K., Capati, G., 2010. Web-based knowledge management system: linking smart metering to the future of urban water planning. *Australian Planner* 47 (2), 66–74.
- Stewart, R. A., Willis, R. M., Panuwatwanich, K., Sahin, O., 2013. Showering behavioural response to alarming visual display monitors: longitudinal mixed method study. *Behaviour & Information Technology* 32 (7), 695–711.
- Strengers, Y., 2011. Negotiating everyday life: The role of energy and water consumption feedback. *Journal of Consumer Culture* 11 (3), 319–338.
- Suero, F., Mayer, P., Rosenberg, D., 2012. Estimating and verifying united states households' potential to conserve water. *Journal of Water Resources Planning and Management* 138 (3), 299–306.
- Suzuki, K., Inagaki, S., Suzuki, T., Nakamura, H., Ito, K., 2008. Nonintrusive appliance load monitoring based on integer programming. In: *SICE Annual Conference, 2008. IEEE*, pp. 2742–2747.
- Syme, G. J., Shao, Q., Po, M., Campbell, E., 2004. Predicting and understanding home garden water use. *Landscape and Urban Planning* 68 (1), 121 – 128.
- Talebpour, M., Sahin, O., Siems, R., Stewart, R., 2014. Water and energy nexus of residential rain water tanks at an end use level: Case of australia. *Energy and Buildings*.
- Thomas, J. F., Syme, G. J., 1988. Estimating residential price elasticity of demand for water: A contingent valuation approach. *Water Resources Research* 24 (11), 1847–1857.
- Tibshirani, R., 1996. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society. Series B* 58 (1), 267–288.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K., 2005. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67 (1), 91–108.
- Tropp, J., 2004. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. on Information Theory* 50 (10), 2231–2242.
- UNDESA, 2010. World urbanization prospects: The 2014 Revision. Highlights (ST/ESA/SER.A/352). United Nations Department of Economic and Social Affairs Population Division.
- Van Der Linden, S., 2013. Exploring beliefs about bottled water and intentions to reduce consumption: The dual-effect of social norm activation and persuasive information. *Environment and Behavior*, 0013916513515239.
- Van Der Maaten, L., 2014. Accelerating t-sne using tree-based algorithms. *Journal of machine learning research* 15 (1), 3221–3245.
- Vassileva, I., Campillo, J., 2014. Increasing energy efficiency in low-income households through targeting awareness and behavioral change. *Renewable Energy* 67, 59 – 63, renewable Energy for Sustainable Development and Decarbonisation.

Bibliography

- Verdú, S. V., Garcia, M. O., Senabre, C., Marín, A. G., Franco, F. J. G., 2006. Classification, filtering, and identification of electrical customer load patterns through the use of self-organizing maps. *IEEE Transactions on Power Systems* 21 (4), 1672–1682.
- Vieira, P., Loureiro, D., Barateiro, J., Ribeiro, R., Rebelo, M., Coelho, S. T., 2014. Innovative smart metering-based applications for water utilities. *Water Utility Management International* 9, 16.
- Vincent, T., Novara, C., 2013. Mixed parametric/non-parametric identification of systems with discontinuous nonlinearities. *Automatica* 49 (12), 3661–3669.
- Viterbi, A. J., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *Information Theory, IEEE Transactions on* 13 (2), 260–269.
- Vogt, J. E., Roth, V., 2010. The group-lasso: $\ell_{1,\infty}$ regularization versus $\ell_{1,2}$ regularization. In: *Pattern Recognition*. Springer, pp. 252–261.
- WaterSmart software, 2014. Tapping into the power of behavioral science - insights & opportunities for water-use efficiency.
- Wentz, E. A., Gober, P., 2007. Determinants of small-area water consumption for the city of phoenix, arizona. *Water Resources Management* 21 (11), 1849–1863.
- Willis, R., Stewart, R. A., Panuwatwanich, K., Capati, B., Giurco, D., 2009b. Gold coast domestic water end use study. *Water* 36 (6), 79–85.
- Willis, R. M., Stewart, R. A., Capati, B., 2009a. Closing the loop on water planning: an integrated smart metering and web-based knowledge management system approach. In: *ICA 2009*. International Water Association.
- Willis, R. M., Stewart, R. A., Giurco, D. P., Telebpour, M. R., Mousavinejad, A., 2013. End use water consumption in households: impact of socio-demographic factors and efficient devices. *Journal of Cleaner Production* 60, 107 – 115, special Volume: Water, Women, Waste, Wisdom and Wealth.
- Willis, R. M., Stewart, R. A., Panuwatwanich, K., Jones, S., Kyriakides, A., 2010. Alarming visual display monitors affecting shower end use water and energy conservation in australian residential households. *Resources, Conservation and Recycling* 54 (12), 1117 – 1127.
- Willis, R. M., Stewart, R. A., Panuwatwanich, K., Williams, P. R., Hollingsworth, A. L., 2011. Quantifying the influence of environmental and water conservation attitudes on household end use water consumption. *Journal of Environmental Management* 92 (8), 1996–2009.
- Wong, J. S., Zhang, Q., Chen, Y. D., 2010. Statistical modeling of daily urban water consumption in hong kong: Trend, changing patterns, and forecast. *Water Resources Research* 46 (3).
- Wooldridge, M., 2009. *An Introduction to MultiAgent Systems*, 2nd Edition. Wiley, NY.
- Worthington, A. C., Hoffman, M., 2008. An empirical survey of residential water demand modelling. *Journal of Economic Surveys* 22 (5), 842–871.
- WWAP, 2014. *The United Nations World Water Development Report 2014: Water and Energy*. Paris, UNESCO.
- Young, R. A., 1973. Price elasticity of demand for municipal water: A case study of tucson, arizona. *Water Resources Research* 9 (4), 1068–1072.
- Yuan, M., Lin, Y., 2006. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B* 68 (1), 49–67.

- Zarli, A., Rezgui, Y., Belziti, D., Duce, E., 2014. Water analytics and intelligent sensing for demand optimised management: The wisdom vision and approach. *Procedia Engineering* 89, 1050–1057.
- Zeifman, M., Akers, C., Roth, K., 2011. Nonintrusive Appliance Load Monitoring (NIALM) for Energy Control in Residential Buildings: Review and Outlook. In: *IEEE transactions on Consumer Electronics*. Citeseer.
- Zhao, B., Stankovic, L., Stankovic, V., 2015. Blind non-intrusive appliance load monitoring using graph-based signal processing. In: *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, pp. 68–72.
- Zhao, B., Stankovic, L., Stankovic, V., 2016. On a Training-Less Solution for Non-Intrusive Appliance Load Monitoring Using Graph Signal Processing. *IEEE Access* 4, 1784–1799.
- Zhao, P., Rocha, G., Yu, B., 2009. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 3468–3497.
- Zhao, Z., Morstatter, F., Sharma, S., Alelyani, S., Anand, A., Liu, H., 2010. Advancing feature selection research. *ASU Feature Selection Repository*.
- Zhou, S. L., McMahon, T. A., Walton, A., Lewis, J., 2000. Forecasting daily urban water demand: a case study of melbourne. *Journal of Hydrology* 236 (3), 153–164.
- Zhou, S. L., McMahon, T. A., Walton, A., Lewis, J., 2002. Forecasting operational demand for an urban water supply zone. *Journal of Hydrology* 259 (1-4), 189–202.
- Zoha, A., Gluhak, A., Imran, M. A., Rajasegarar, S., 2012. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors* 12 (12), 16838–16866.
- Zou, H., 2006. The adaptive lasso and its oracle properties. *Journal of the American statistical association* 101 (476), 1418–1429.

