

DISS. ETH NO. 28039

# Marrying Machine Learning and Model Predictive Control for efficient Building Energy Management

A thesis submitted to attain the degree of  
DOCTOR OF SCIENCES of ETH ZURICH  
(Dr. sc. ETH Zurich)

presented by

FELIX BÜNNING

M. Sc. in Energy Engineering, RWTH Aachen University  
born on 08. September 1990  
citizen of Germany

accepted on the recommendation of

Prof. Dr. John Lygeros, examiner (ETH Zurich, Switzerland)  
Philipp Heer, co-examiner (Empa, Switzerland)  
Prof. Dr. Les Norford, co-examiner (MIT, USA)

2021

---

© December 2021  
Felix Bünning  
All Rights Reserved

ISBN 978-3-907234-89-1  
DOI 10.3929/ethz-b-000526883



To my friends and my family.

---

*Und jedem Anfang wohnt ein Zauber inne,  
Der uns beschützt und der uns hilft, zu leben.*

Stufen, Hermann Hesse (1877–1962)

# Abstract

The heating and cooling of buildings are responsible for a significant share of the global final energy consumption and CO<sub>2</sub> emissions. At the same time, buildings and districts with electric heating and cooling systems can facilitate the integration of volatile renewable energy sources in the electricity grid by providing reserves on the demand side. By shifting their demand in time depending on the availability of renewable electricity, they can contribute to stabilizing the grid. A suitable control approach to minimize the energy consumption of buildings and districts, and to provide electric reserves with them is Model Predictive Control (MPC), where control inputs are periodically optimized over a prediction horizon, given a model of the system dynamics, measurements, disturbance forecasts, constraints, and a cost function. However, as buildings are individual, developing and maintaining models based on physics for building dynamics and demand forecasting, which serve as inputs to the controllers, are often considered tedious and as a result cost-prohibitive for industrial application. This thesis, therefore, aims to find scalable, data-driven methods that simplify the modeling process, and which combined with the concept of predictive control, yield viable controllers for practical building energy and district energy applications.

First, we present online correction methods for forecasting the heating demand of buildings with Artificial Neural Networks (ANN), based on the autocorrelation of the forecasting error and online learning. We show that these significantly increase the forecasting accuracy in case studies based on measurement data of four buildings. Moreover, they reduce the variance between the forecasts of different networks trained on the same data, caused by the non-convex training process of ANN. Furthermore, two Machine Learning methods, based on Random Forests and Input Convex Neural Networks, for data-driven modeling of building thermal dynamics are presented, as well as a method based on physics-informed Autoregressive Moving Average with Extra Input (ARMAX) models. The methods, which lead to convex MPC formulations, are compared in a series of experiments in an occupied apartment. Here, it is found that data-driven MPC reduces the heating and cooling energy consumption by 26% to 49% compared to an industry-standard hysteresis controller in the given case. The physics-informed ARMAX model shows better accuracy and sample efficiency compared to the Machine Learning methods.

These modeling methods are then combined with a Robust MPC scheme to provide day-ahead electrical reserves with a district system comprising a central heat pump and buffer storage. Here, we distinguish between connected buildings that participate in the

---

reserve scheme, which means that their thermal inertia can be exploited by the central controller, and buildings that do not participate due to privacy concerns or hardware limitations. The dynamics of the former are modeled with the presented ARMAX method, whereas the thermal demands of the latter are treated as an aggregated disturbance that is forecast with the presented data-driven demand forecasting method. Only industrial components are modeled with first principles in the MPC scheme. We use affine policies on uncertain disturbances such as the regulation signal to mitigate the effect of uncertainty accumulation in the controlled state. As a result, the reserves offered are increased significantly. We demonstrate that even in a system configuration where all connected buildings do not actively participate in the scheme, considerable amounts of reserves can be offered. However, we also show that by exploiting the thermal inertia of a subset of the connected buildings, the reserves offered can be increased by a factor of 2.7. Both cases, with buildings being actively controlled and not being actively controlled, are validated in experiments, where it is demonstrated that the comfort in the connected building is ensured, while reserves are being provided and the heat pump is tracking a regulation signal with high accuracy.

In conclusion, it is experimentally demonstrated that data-driven methods can lead to viable, reliable and scalable controllers that facilitate energy savings and electrical reserve provision in buildings and districts. We find that physics-based priors enforced in a linear regression process, can lead to higher model accuracy compared to more flexible but system-agnostic Machine Learning methods, even if abundant training data is available. By combining data-driven methods with Robust MPC using affine policies and physics-based models of some industrial components, even complex control tasks, such as day-ahead reserves provision in buildings and districts can be addressed.

# Zusammenfassung

Das Heizen und Kühlen von Gebäuden ist für einen erheblichen Anteil des globalen Endenergieverbrauchs und der CO<sub>2</sub>-Emissionen verantwortlich. Gleichzeitig können Gebäude und Quartiere mit elektrischen Heiz- und Kühlsystemen die Integration von volatilen erneuerbaren Energiequellen in das Stromnetz erleichtern, indem sie auf der Verbraucherseite elektrische Reserven bereitstellen. Indem sie ihre Nachfrage in Abhängigkeit von der Verfügbarkeit erneuerbarer Energie zeitlich verschieben, können sie zur Stabilisierung des Netzes beitragen. Ein geeigneter Regelansatz zur Minimierung des Energieverbrauchs von Gebäuden und Quartieren und zur Bereitstellung elektrischer Reserven ist die modellprädiktive Regelung (engl. Model Predictive Control, MPC). Hier werden die Stellgrößen periodisch bezüglich eines Vorhersagehorizonts, eines Modells der Systemdynamik, bezüglich Messungen, Vorhersagen von Störgrößen, Bedingungen und einer Kostenfunktion optimiert. Da jedes Gebäude individuell ist, ist die Entwicklung und Erhaltung von physikalischen Modellen, sowohl für die Gebäudedynamik als auch für Wärmebedarfsvorhersagen, welche häufig als Störgrößen auf höheren Regelebenen dienen, jedoch oft mühsam und deshalb für kommerzielle Anwendungen unrentabel. Ziel dieser Arbeit ist es daher, skalierbare datengetriebene Methoden zu finden, die den Modellierungsprozess vereinfachen und in Kombination mit dem Konzept der prädiktiven Regelung praktikable Regler für Gebäude- und Quartierenergie Anwendungen liefern.

Zunächst stellen wir Online-Korrekturmethode für die Vorhersage des Wärmebedarfs von Gebäuden mit Künstlichen Neuronalen Netzen (ANN) vor, die auf der Autokorrelation des Vorhersagefehlers und Online-Learning basieren. Wir zeigen, dass diese Methoden in Fallstudien, welche auf Messdaten von vier Gebäuden basieren, die Vorhersagegenauigkeit deutlich erhöhen. Außerdem verringern sie die Varianz zwischen den Prognosen verschiedener Netzwerke, die auf denselben Daten trainiert wurden - ein Effekt, der auf den nicht-konvexen Trainingsprozess von ANN zurückzuführen ist. Darüber hinaus werden zwei Methoden des maschinellen Lernens, auf der Grundlage von Random Forests und Input Convex Neural Networks, für die datengetriebene Modellierung der thermischen Dynamik von Gebäuden, sowie eine Methode auf der Grundlage von Physik-inspirierten Autoregressive Moving Average with Extra Input (ARMAX) Modellen vorgestellt. Alle Methoden führen zu konvexen MPC-Formulierungen und werden in einer Reihe von Experimenten in einem bewohnten Apartment getestet und verglichen. Dabei wird gezeigt, dass die datengetriebenen MPC Ansätze den Heiz- und Kühlenergieverbrauch im Vergleich zu einem Standard-Hystereseregler im vorliegenden Fall um 26%

---

bis 49% reduzieren. Das Physik-inspirierte ARMAX Modell zeigt im Vergleich zu den Methoden des maschinellen Lernens eine bessere Genauigkeit und Dateneffizienz.

Diese Modellierungsmethoden werden dann mit einem robusten MPC-Schema kombiniert, um mit einem Fernwärmesystem mit zentraler Wärmepumpe und einen Pufferspeicher elektrische Day-Ahead-Reserven bereitzustellen. Bei den angeschlossenen Gebäuden unterscheiden wir zwischen jenen, die am aktiv am Regelsystem zur Bereitstellung von Reserven teilnehmen, was bedeutet, dass ihre thermische Trägheit vom zentralen Regler ausgenutzt werden kann, und anderen, die aufgrund von Datenschutzbedenken oder Hardwarebeschränkungen nicht teilnehmen. Die Dynamik ersterer wird mit der zuvor eingeführten ARMAX-Methode modelliert, während der Wärmebedarf letzterer als eine aggregierte Störgrösse behandelt wird, die mit der vorgestellten datengetriebenen Methode zur Wärmebedarfsvohersage prognostiziert wird. Ausschliesslich industrielle Komponenten werden in dem Schema auf Basis von White-Box Modellen modelliert. Wir nutzen affine Policies um den Effekt von kummulierter Unsicherheit bezüglich der Zustandsvariablen aufgrund von stochastischen Störgrössen (wie zum Beispiel das Regelsignal des Elektrizitätsanbieters zur Abfrage der Reserven) zu verringern. Dadurch können die angebotenen Reserven deutlich erhöht werden. Wir zeigen, dass selbst in einer Systemkonfiguration, in der keines der angeschlossenen Gebäude seine thermische Trägheit zur Verfügung stellt, beträchtliche Mengen an Reserven angeboten werden können. Wir zeigen aber auch, dass durch Ausnutzung der thermischen Trägheit einer Untergruppe der angeschlossenen Gebäude die angebotenen Reserven um den Faktor 2,7 erhöht werden können. Beide Fälle werden in Experimenten validiert, in denen gezeigt wird, dass der Komfort in allen Gebäuden stets gewährleistet ist, während Reserven bereitgestellt werden und die Wärmepumpe einem Regelsignal mit hoher Genauigkeit folgt.

Zusammenfassend wird experimentell nachgewiesen, dass datengetriebene Methoden zu praktikablen, zuverlässigen und skalierbaren Reglern führen können, welche Energieeinsparungen und die Bereitstellung von elektrischen Reserven in Gebäuden und Quartieren ermöglichen. Wir stellen fest, dass wenn physikalische Bedingungen in einem linearen Regressionsprozess berücksichtigt werden, eine höhere Modellgenauigkeit erreicht werden kann im Vergleich zu flexibleren, aber systemagnostischen Methoden des maschinellen Lernens. Dies ist selbst der Fall, wenn eine grosse Menge an Trainingsdaten zur Verfügung steht. Durch die Kombination datengetriebener Modellierungsmethoden mit einem robusten MPC-Schema und affine Policies, sowie White-Box Modellen einiger weniger industrieller Komponenten können selbst komplexe Regelaufgaben wie die Bereitstellung von Day-Ahead-Reserven in Gebäuden und Quartieren gelöst werden.



# Acknowledgements

My first name, Felix, means *lucky* or *fortunate*. Indeed, I owe the possibility to spend four wonderful years at ETH Zürich and Empa and finish them with this thesis to the circumstance that I was fortunate enough to have met many people that helped, guided and accompanied me on the path towards my doctorate. These people are colleagues, mentors, friends and my family.

First and foremost, I would like to thank Prof. John Lygeros for his supervision. Although I had no extensive formal background in control and automation, John gave me the chance to pursue my research at the Automatic Control Laboratory (IfA) of ETH Zürich and his trust to explore different research areas. His leadership style, which always finds the right balance between granting freedom and offering guidance, allowed me to be productive without pressure and challenged me in the right moments. I would also like to thank Philipp Heer, my supervisor at the Urban Energy Systems Laboratory at Empa, for trusting me in setting my own research goals and giving me valuable feedback towards achieving them. I am also grateful to Prof. Roy Smith's guidance, especially in theoretical questions, from which I benefited greatly. Special thanks also go to Dr. Kristina Orehounig and Viktor Dorer for giving me the opportunity to conduct my research at the Urban Energy Systems Laboratory. Moreover, I would like to thank Prof. Les Norford at MIT for agreeing to be on my examination committee and providing valuable and encouraging feedback for this thesis.

I was also lucky to work with many kind, smart and helpful colleagues. I would specifically like to thank Dr. Georgios Mavromatidis, Dr. Christoph Waibel, Dr. Marc Hohmann, Dr. Somil Miglani, Dr. Portia Murray, Dr. Cristina Dominguez, Dr. Dan-hong Wang and Dr. Julien Marquant for giving me such a warm welcome to the Empa group. I would also like to thank my other Empa colleagues for providing such a great working atmosphere, especially Dr. James Allan, Dr. Natasa Vulic, Dr. Mashaël Yazdanie, Dr. Hanmin Cai, Varsha Behrunani and Loris Di Natale. At IfA, I would especially like to thank Dr. Benjamin Flamm, Dr. Mohammad Khosravi, Dr. Joseph Warrington, Dr. Mathias Hudoba de Badyn, Dr. Annika Eichler, Ahmed Aboudonia, Marta Fochesato, Francesco Micheli, Dr. Andrea Iannelli and Anil Parsi for many fruitful discussions in the building energy group and their excitement to listen to my applied research talks.

During my research, I collaborated with many people from whom I benefited greatly. Here, I would like to specifically thank Ahmed Aboudonia for successfully co-supervising

---

many student theses with me. Special thanks also go to Benjamin Huber, my first Master student and now trusted colleague, whose work has greatly contributed to my research. Moreover, I am grateful to Sascha Stoller and Reto Fricker for providing and maintaining the NEST infrastructure, which was essential for my experiments. I would also like to thank Dr. Joseph Warrington, Dr. Mathias Hudoba de Badyn, Dr. Georgios Mavromatidis and Dr. Hanmin Cai for valuable feedback on my research articles and this thesis.

Along my way towards starting a doctorate, I was supported and mentored by many inspiring people, starting with Prof. Piero Colonna at TU Delft, who sparked my interest in the dynamic simulation of energy systems. I would specifically like to thank Dr. Marcus Fuchs, Prof. Dirk Müller and Doris Heinzl of the Institute for Energy Efficient Buildings and Indoor Climate at RWTH Aachen for the opportunity and guidance to publish as a young researcher, and for their efforts to support my research visit to Berkeley Lab. I am deeply thankful to Dr. Michael Wetter at Berkeley Lab for his encouragement to pursue a PhD, his ongoing interest in my research and for being a role model in terms of both rigor and efficiency in conducting research.

People that have worked closely with me know, that my ideas commonly are not formed at the desk under working pressure but often come to me rather spontaneously and are a direct result of me being relaxed and feeling content. I owe this entire circumstance of contentment to the people surrounding me, which are my friends and family. I would, therefore, like to thank Andreas Behler and Thilo Wagenführ, with whom I have spent six great years in Aachen as flatmates and friends. Special thanks also go to Eliza Girod and Rik Bähnemann, who have accompanied me during my whole PhD time in Zürich, from hosting me on my interview day, over uncounted times chatting, hiking, having dinner, etc., all the way to helping me with the PhD Aperó. I would also like to thank Julian Merkel, Claudia Wiehler and all other Zürich friends for their support and all the fun times we had together. I am also thankful to Dr. Augusto Gandia, who made my start in Zürich so easy and fun.

I am forever thankful to Dr. Kathrin Wenzel for her endless support on my path, including her encouragement and acceptance for me to spend time abroad, her resulting adaptations to her own personal life, and even more importantly, for her witty, interested, and caring nature, and for her love that keeps me happy, content and rooted. Lastly, I want to thank my family, my parents Dorothee Dosquet and Hartmut Bünning, and my brothers Niklas and Hendrik Bünning, for providing the foundation of all of my endeavors, for supporting my education, for their love and trust, and their ongoing encouragement. I am truly blessed to have such great friends and family.

Felix Bünning  
Zürich, 2021

# List of Figures

1.1	Overview on the topics discussed in this thesis. . . . .	5
2.1	Scheme of forecast correction based on autocorrelation of forecasting errors	14
2.2	NEST building at Empa in Switzerland . . . . .	16
2.3	NEST heating system schematic . . . . .	17
2.4	Measured total heating load of the NEST building. . . . .	18
2.5	Measured ambient temperature at the roof of NEST building. . . . .	18
2.6	RC building models with varying numbers of resistors and capacitors. . .	21
2.7	Coefficient of determination for ANN and forecast correction methods for 100 individual networks. . . . .	24
2.8	Coefficient of determination of 2h ahead prediction of 100 ANN and fore- cast correction methods. . . . .	25
2.9	Coefficient of determination of 24h ahead prediction of 100 ANN and forecast correction methods. . . . .	26
2.10	Coefficient of determination of 100 ANN and forecast correction methods with randomized training period. . . . .	27
2.11	Autocorrelation of the residual uncorrected ANN, and ANN with correc- tion by error-autocorrelation and online learning. . . . .	28
2.12	Forecasting error histogram for uncorrected ANN, and ANN with correc- tion by error-autocorrelation and online learning. . . . .	29
2.13	Forecasting examples compared to the real heating load with varying tem- poral resolution. . . . .	30
2.14	Buildings used in the validation study. . . . .	34
2.15	Variance of the coefficient of determination for all tested buildings. . . . .	35
2.16	Forecasting trajectory example from building (b) test set. . . . .	37
2.17	Variance of the coefficient of determination for different sampling times. .	39
2.18	Variance of the coefficient of determination for different noise levels. . . .	40
2.19	Forecast trajectory for test set with both correction methods applied with single week of offline training. . . . .	41
2.20	Variance of the coefficient of determination for different uses of weather forecasts with uncorrected demand forecasts, and corrected forecasts. . .	42

3.1	Schematic of the system under consideration with heat pump, water storage tank and heat exchangers for individual apartments. . . . .	50
3.2	Hierarchical control scheme. . . . .	51
3.3	Forecast correction based on error-autocorrelation and online learning. . . . .	53
3.4	Experimental facility with heat pump and water storage tanks. . . . .	60
3.5	Experimental results of day 1. . . . .	63
3.6	Experimental results of day 2. . . . .	64
3.7	Experimental results of day 3. . . . .	65
3.8	Results tracking performance. . . . .	67
3.9	Reserves offered in Level 1: affine policies vs. open-loop MPC. . . . .	70
3.10	Cost in Level 1: affine policies vs. open-loop MPC. . . . .	70
3.11	Reserves offered in Level 1: affine policies vs. perfect knowledge. . . . .	71
3.12	Cost in Level 1: affine policies vs. perfect knowledge. . . . .	72
3.13	Cost in Level 1: affine policies vs. perfect knowledge, in detail. . . . .	72
4.1	Schematic of a Fully Input Convex Neural Network . . . . .	83
4.2	Schematic of a Partially Input Convex Neural Network . . . . .	83
4.3	Rendering of the UMAR unit in NEST with both bedrooms. . . . .	91
4.4	Heating system of the UMAR unit in NEST. . . . .	91
4.5	Cooling experiment with ARMAX in bedroom 1 and FICNN in bedroom 2. . . . .	96
4.6	Cooling experiment with RF in bedroom 1 and bedroom 2. . . . .	97
4.7	Heating energy of MPC methods and baseline controller as a function of HDS and CDS. . . . .	100
4.8	Mean Squared Error for 1-hour open loop prediction with ICNN, RF and ARMAX models. . . . .	102
4.9	Mean Squared Error for 1-hour open loop prediction with ARMAX model and varying control inputs. . . . .	103
5.1	Schematic of the system under consideration. . . . .	110
5.2	Hierarchical control scheme with two control Levels. . . . .	111
5.3	DFAB HOUSE at NEST at Empa in Switzerland. . . . .	120
5.4	Results of the numerical case study. . . . .	123
5.5	Results of the control experiment at NEST. . . . .	126
5.6	Room temperature examples from the units not directly controlled by the controller. . . . .	127
5.7	Detailed plot of the heat pump following the regulation signal. . . . .	127
A.1	Mean Squared Error of 100 ANN and forecast correction methods. . . . .	135
A.2	Mean Absolute Error of 100 ANN and forecast correction methods. . . . .	136

A.3	Coefficient of Variation of Root-Mean Squared Error of 100 ANN and forecast correction methods. . . . .	136
A.4	Scaled forecasting trajectories on the training set for buildings (a)-(d). . .	137
A.5	Scaled forecasting trajectories on the test set for buildings (a)-(d). . . .	138
A.6	Example from the validation set of the one-hot solar model. . . . .	139
A.7	Baseline controller example for the heating case in bedroom 2. . . . .	139
A.8	Temperature analysis of MPC vs. baseline controller for heating experiments in bedroom 1. . . . .	140

*List of Figures*

---

## List of Tables

2.1	ANN parameters . . . . .	20
2.2	Coefficient of determination for all tested heating demand forecasting methods. . . . .	31
2.3	$R^2$ of different modeling approaches for all tested buildings. . . . .	36
3.1	Parameters for controller Levels 1 and 2. . . . .	60
4.1	Computational resources for MPC controllers with varying models. . . . .	102
5.1	Parameters for controller Level 1 related to heat pump and storage. . . . .	121





# Contents

<b>Abstract</b>	<b>i</b>
<b>Zusammenfassung</b>	<b>iii</b>
<b>Acknowledgements</b>	<b>v</b>
<b>List of Figures</b>	<b>vii</b>
<b>List of Tables</b>	<b>ix</b>
<b>Contents</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.1.1 Demand prediction in buildings and districts . . . . .	2
1.1.2 Data-driven modeling for predictive building control . . . . .	3
1.1.3 Reserve provision with buildings and districts . . . . .	3
1.2 Outline and contributions . . . . .	4
1.3 Publications contributing to the thesis . . . . .	6
1.4 Other publications . . . . .	8
<b>2 Demand Forecasting with online-corrected ANN</b>	<b>9</b>
2.1 Introduction . . . . .	9
2.1.1 Contribution . . . . .	11
2.1.2 Structure . . . . .	12
2.2 Problem statement . . . . .	13
2.3 Methodology . . . . .	13
2.3.1 Error-autocorrelation correction . . . . .	13
2.3.2 Online learning . . . . .	14
2.4 Case study 1: General Validation . . . . .	16
2.4.1 System description . . . . .	16
2.4.2 Artificial Neural Network Configuration . . . . .	17
2.4.3 Forecasting methods for benchmarking . . . . .	20

2.4.4	Results and discussion . . . . .	23
2.5	Case study 2: Transferrability to other Buildings . . . . .	32
2.5.1	General Description . . . . .	33
2.5.2	Results and discussion . . . . .	33
2.6	Case study 3: Sensitivity with respect to data quality and quantity . . .	37
2.6.1	General description . . . . .	38
2.6.2	Results and Discussion . . . . .	39
2.7	Conclusion . . . . .	42
<b>3</b>	<b>Robust MPC for frequency regulation with heat pumps</b>	<b>45</b>
3.1	Introduction . . . . .	46
3.1.1	Contribution . . . . .	48
3.1.2	Structure . . . . .	49
3.2	Problem statement . . . . .	49
3.2.1	Reserve provision scheme . . . . .	49
3.2.2	System under consideration . . . . .	50
3.2.3	Control scheme . . . . .	51
3.3	Methodology . . . . .	51
3.3.1	Models . . . . .	51
3.3.2	Controller design . . . . .	54
3.4	Experimental case study . . . . .	59
3.4.1	Configuration . . . . .	59
3.4.2	Results . . . . .	62
3.5	Numerical case study . . . . .	68
3.5.1	Use of affine policies compared to open-loop MPC . . . . .	69
3.5.2	Use of affine policies compared to a system with perfect knowledge	70
3.6	Limitations and future directions . . . . .	73
3.7	Conclusion . . . . .	74
<b>4</b>	<b>Data-driven predictive building control</b>	<b>75</b>
4.1	Introduction . . . . .	75
4.1.1	Contribution . . . . .	78
4.1.2	Structure . . . . .	79
4.2	Problem statement . . . . .	79
4.3	Methodology . . . . .	80
4.3.1	Random Forests with linear regression leaves . . . . .	80
4.3.2	Input convex neural networks . . . . .	82
4.3.3	Physics-informed ARMAX models . . . . .	86

4.3.4	Model properties . . . . .	90
4.4	Test bed . . . . .	91
4.5	Experiment case study . . . . .	92
4.5.1	Model and controller configuration . . . . .	93
4.5.2	Example closed-loop experiments . . . . .	95
4.5.3	Energy consumption . . . . .	97
4.5.4	Computational performance . . . . .	101
4.6	Numerical case study . . . . .	102
4.7	Conclusion . . . . .	104
<b>5</b>	<b>Combining electric reserves with data-driven building control</b>	<b>107</b>
5.1	Introduction . . . . .	107
5.1.1	Contribution . . . . .	108
5.1.2	Structure . . . . .	109
5.2	Problem statement . . . . .	110
5.2.1	Reserve provision scheme . . . . .	110
5.2.2	System under consideration . . . . .	110
5.2.3	Control scheme . . . . .	111
5.3	Methodology . . . . .	112
5.3.1	Modeling . . . . .	112
5.3.2	Controller design . . . . .	115
5.4	Numerical case study and results . . . . .	120
5.4.1	Configuration . . . . .	120
5.4.2	Results and discussion . . . . .	122
5.5	Experimental case study and results . . . . .	124
5.5.1	Configuration . . . . .	125
5.5.2	Results and discussion . . . . .	125
5.6	Limitations . . . . .	129
5.7	Conclusion . . . . .	130
<b>6</b>	<b>Conclusions and future directions</b>	<b>131</b>
6.1	Summary and conclusion . . . . .	131
6.2	Future directions . . . . .	133
<b>A</b>	<b>Appendix</b>	<b>135</b>
A.1	Additional KPI for Case study 1 in Chapter 2 . . . . .	135
A.2	Forecasting trajectories for Case study 2 in Chapter 2 . . . . .	137
A.3	Additional Figures for Chapter 4 . . . . .	139

*Contents*

---

<b>Bibliography</b>	<b>141</b>
<b>Curriculum Vitae</b>	<b>157</b>

# Introduction

## 1.1 Motivation

The operation of buildings is responsible for 28% of the global energy-related CO<sub>2</sub> emissions (Abergel et al., 2019), of which a large fraction is caused by space heating and cooling (Drgoňa et al., 2020a). In addition to retrofitting heating and cooling equipment and the envelope of a building, advanced control methods can be used to reduce the energy consumption of individual buildings (Široký et al., 2011; Yudong Ma et al., 2012; Oldewurtel et al., 2010; Sturzenegger et al., 2016; Hameed Shaikh et al., 2014), district heating and cooling systems (Bünning et al., 2018; Hohmann et al., 2019; van der Heijde et al., 2019), and energy hubs (Geidl et al., 2007; Arnold et al., 2009; Arnold and Andersson, 2011; Darivianakis et al., 2017).

While the operation of buildings contributes to climate change through direct and indirect CO<sub>2</sub> emissions, at the same time buildings can increase the potential to integrate renewable energy sources into the electricity grid (Langevin et al., 2021). As many of these sources are highly volatile, there is a growing need for frequency regulation (Johnson et al., 2019). A common strategy for frequency regulation is the deployment of fast-reacting power plants, for example gas or hydro-power; an emerging strategy is the use of storage technologies. However, many storage technologies, such as batteries, are still relatively cost intensive, and fossil fuel-based power plants will be phased out in the future. Apart from such regulation on the supply side of the grid, frequency regulation on the demand side is possible through the manipulation of controllable loads. This concept falls under the category of demand-side management (Palensky and Dietrich, 2011; Gelazanskas and Gamage, 2014). Buildings equipped with electric heating or cooling systems are possible candidates for demand-side management (Kim et al., 2016; Blum et al., 2016). Due to their thermal inertia and the insensitivity of humans to minor temperature fluctuations, buildings are flexible to some extent when it comes to their heating and cooling requirements, hence their electricity consumption. By shifting their consumption in time, they can therefore influence the grid frequency (Fischer and Madani, 2017).

In both applications, i.e. reducing the energy consumption of buildings, and using buildings to provide services to the electricity grid, the comfort of the occupants needs to be ensured. For this, many control approaches in the literature (Zhang et al., 2017; Vrettos et al., 2016; Oldewurtel et al., 2012; Sturzenegger et al., 2016; Drgoňa et al., 2020a; Henze, 2013; Hilliard et al., 2017) use Model Predictive Control (MPC) (Morari and H. Lee, 1999; Kouvaritakis and Cannon, 2016), which allows to calculate optimal control inputs over a receding horizon, given a cost function, a set of constraints, disturbance measurements and forecasts, and a model of the dynamics of the considered system. However, models in building MPC are conventionally based on physics (Picard et al., 2015; Sturzenegger et al., 2014), which means that the model is built using principles of heat transfer and thermodynamics, or based on expensive excitation experiments, or a combination of both. As buildings differ from each other and thus need to be modeled individually, manually developing and maintaining such models is often considered too expensive to justify investment, an issue that potentially hinders the commercial application of MPC in buildings (Sturzenegger et al., 2016). Similar issues arise for first-principles-based heating and cooling demand forecasting methods (Wang and Xu, 2006), which serve as disturbance forecasts for MPC on the district level.

As in many other domains (Silver et al., 2017; Wojna et al., 2018; Tobin et al., 2017), Machine Learning based or data-driven methods have also received increasing attention in the area of building control. Their application includes heating and cooling demand forecasting (Mat Daut et al., 2017; Zhao and Magoulès, 2012; Harish and Kumar, 2016; Wang and Srinivasan, 2017; Foucquier et al., 2013; Suganthi and Samuel, 2012; Amasyali and El-Gohary, 2018; Ahmad et al., 2018), reinforcement learning for building energy management (Mocanu et al., 2019; Pinto et al., 2021; Wang and Hong, 2020), and the generation of dynamical models for building MPC (Wang et al., 2019; Chen et al., 2019; Mugnini et al., 2020; Smarra et al., 2018). Similar to grey-box modeling approaches, but often with greater flexibility on the parameterization, data-driven methods allow to generate building models directly from measurement data - thus potentially decreasing the cost of development and maintenance of predictive control and prediction methods in buildings.

In this thesis, we aim to leverage data-driven methods to address the challenges of demand forecasting, predictive building control, and electric reserve provision with buildings and districts. Here, several gaps in literature are addressed, which are briefly discussed for each individual field in the following. More detailed backgrounds are given in the individual chapters.

### 1.1.1 Demand prediction in buildings and districts

Accurate heating and cooling demand predictions are essential inputs to predictive controllers addressing problems such as energy-hub operation optimization (Darivianakis

et al., 2017) or electric reserve provision with buildings and districts, as presented in this thesis. With increasing availability of demand measurement data, data-driven forecasting methods have received growing attention in the literature, for individual buildings (Paudel et al., 2014; Mestekemper et al., 2013) and districts (Suryanarayana et al., 2018; Saloux and Candanedo, 2018). Many of these studies report promising results, but none of them communicates whether the achieved performance are average results or best case results. This is problematic as Machine Learning methods often suffer from high variance due to the stochastic nature of the initialization and training processes. Ensemble methods (Jovanović et al., 2015; Jetcheva et al., 2014) address the variance issue, but through averaging can lead to sub-optimal prediction accuracy. More importantly they are computationally expensive, which limits their suitability for the use in online control applications.

### 1.1.2 Data-driven modeling for predictive building control

Obtaining first-principles based models for predictive building control is expensive because each building is individual. As a result, data-driven building modeling approaches that rely purely on historical measurement data have emerged (Smarra et al., 2018). While many of these methods applied to predictive control have been tested in simulation (Mugnini et al., 2020; Wang et al., 2019), the application and validation on real systems have not been appropriately addressed in the literature so far (Kathirgamanathan et al., 2021). Methods that are indeed validated in practical applications of predictive building control (Jain et al., 2020; Yang et al., 2020) often rely on models that lead to non-convex problem formulations. While these might be suitable for small control problems such as single-building energy management, they can lead to computationally intractable optimization problems when applied to more complex control tasks, such as distributed building control (Lefebure et al., 2021), or electric reserve provision. Moreover, in the spirit of Mania et al. (2018), it is still unclear how Machine-Learning based models compare to simpler model identification methods in practical applications.

### 1.1.3 Reserve provision with buildings and districts

Several authors (Zhang et al., 2017; Vrettos et al., 2016) have developed robust or stochastic MPC schemes to provide day-ahead reserves for frequency regulation with commercial buildings and HVAC systems. Some have also applied them in practice (Vrettos et al., 2018a). These schemes ensure occupant comfort in the face of uncertainty, introduced for example through the stochastic regulation signal from the transmission system operator. As building energy systems and related components (such as storage tanks) have integrator dynamics, uncertain disturbances lead to built-up uncertainty in the system state and thus limit the electrical reserves offered. This is the case because much of

the available demand flexibility is used by accounting for the disturbance uncertainty. While feedback policies are a known tool in stochastic and Robust MPC to mitigate this problem (Goulart et al., 2006), they have not been applied in practice to reserve provision with buildings and districts yet. Furthermore, so far practical implementations of reserve provision schemes are limited to individual buildings, which have to be heavily aggregated in order to meet minimal capacity requirements for reserve products. As a result, there is a growing interest in the potential of providing ancillary services in district heating systems (Lund et al., 2018; Xu et al., 2020; Pagh Nielsen and Sørensen, 2021; Ivanova et al., 2019). If operated with central heat pumps, they could meet the capacity requirements without the need of aggregation. Review studies (Vandermeulen et al., 2018; Hennessy et al., 2019), however, indicate that there is little work on the operation (Terreros et al., 2020; Salpakari et al., 2016; Li et al., 2016) of such systems and no implementation and validation on a real system in the literature. Moreover, to efficiently design predictive control schemes for such systems, dynamic building models are necessary to model the behavior of households that actively participate in the scheme by offering the thermal inertia of their buildings for reserve provision. Following the reasoning in the previous section, these models should preferably be data-driven and should lead to convex MPC formulations. Equipment such as the central heat pump itself and central storage can be modeled based on physics, as these are industrial products and only need to be modeled once. Other households might refuse to actively participate, for example because of privacy or data security concerns, or might not be suitable to be actively controlled by the central controller due to missing communication infrastructure or hardware. For these buildings, accurate data-driven demand forecasting methods are essential, as also the uncertainty related to these forecasts limits reserve provision.

## 1.2 Outline and contributions

Following the motivation given in the previous section, we present new methods in the areas of online-corrected demand forecasting and data-driven dynamic building models that lead to convex predictive control problem formulations. We combine them with physics-based models of industrial components, and robust control methods to address electrical reserve provision with buildings and districts. They are applied to and validated on real building and district energy systems to address the issues of reducing building energy consumption and demand-side frequency regulation, and comparisons are made between various methods. Specifically, as shown in Figure 1.1, we make the following contributions and organize the thesis as follows:

In Chapter 2, we develop two correction methods which are applied online for heating demand forecasting with Artificial Neural Networks (ANN), to obtain reliable forecasting models. The first one is based on the forecasting error-autocorrelation, and the second



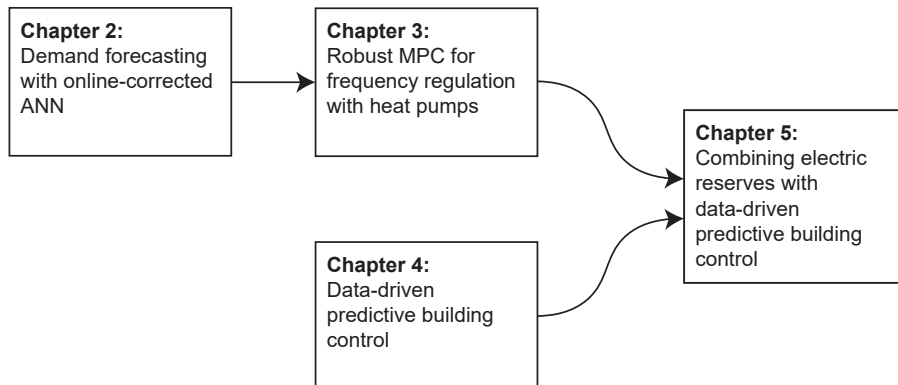


Figure 1.1: Overview on the topics discussed in this thesis.

one is based on online learning. The approach is tested in three case studies in the frame of day-ahead sub-hourly forecasting. In the first case study of a complex building, which has properties of a district heating system, it is demonstrated that the methods significantly reduce variance in prediction performance and also increase average prediction accuracy. When compared to other grey-box and black-box forecasting models, the approach performs well. In the second study, the approach is validated on three additional buildings where the performance of the approach is confirmed. In a third case study, the sensitivity with regards to quality and quantity of training data is investigated, and it is shown that the forecasting method only requires little amounts of training data.

In Chapter 3, we present a control scheme to offer frequency reserves with an electrified district heating system. Instead of exploiting the thermal inertia of connected buildings, we use the inertia of buffer storage for providing reserves, reducing the building models to demand forecasts. This configuration resembles the situation where connected buildings in a district system cannot be controlled actively by a central controller, for example because they do not want to participate in a reserve scheme due to privacy concerns, or because of missing communication infrastructure. The presented control scheme paves the way for Chapter 5, where the thermal inertia of a subset of connected buildings is exploited too. By combining robust Model Predictive Control with affine policies, and heating demand forecasting based on Artificial Neural Networks with online correction methods, as described in Chapter 2, we offer frequency regulation reserves and maintain user comfort with a system comprising a heat pump and a storage tank. The robust formulation ensures sufficient storage temperatures in the face of uncertainty, and the affine policies reduce the effect of the uncertainty. In a three-day experiment with a real district-like building energy system, we demonstrate that the scheme is able to offer reserves in a variety of conditions and track a regulation signal while meeting the heating demand of the connected buildings. In additional numerical studies, we demonstrate that using affine policies significantly decreases the cost function and increases the

amount of offered reserves, and we investigate the suboptimality in comparison to an omniscient control system.

In Chapter 4, we explore various methods of data-driven model generation for convex MPC in thermal building control. We introduce Machine Learning models based on Random Forests and Input Convex Neural Networks, as well as physics-informed ARMAX models, and compare them in experiments on a practical building application, and in a numerical case study. We demonstrate that Predictive Control in general leads to significant savings of heating and cooling energy, compared to the baseline hysteresis controller of the building. Moreover, we show that all different model types lead to satisfactory control performance in terms of constraint satisfaction and energy reduction. However, we also see that the physics-informed ARMAX models have a lower computational burden, and a superior sample efficiency compared to the Machine Learning based models. Moreover, even if abundant training data is available, the ARMAX models have a significantly lower prediction error than the Machine Learning models, which indicates that the encoded physics-based prior of the former cannot independently be found by the latter.

In Chapter 5, we combine the methods of Chapters 2-4 and present a two-level control scheme based on Robust MPC with affine policies to offer frequency reserves with a district system, comprising a central heat pump, as well as cold and warm buffer storage tanks. In contrast to Chapter 3, a subset of the connected buildings is considered to be controlled by the central controller, and others are not. For the controlled buildings, we leverage the data-driven modeling methods of Chapter 4 to overcome the problem of physics-based building modeling being tedious as each building is individual. The energy demand of the other buildings is modeled with the methods developed in Chapter 2. In a numerical case study based on one-year historical data of a real system, we compare the approach to a situation where only the buffer storage is used for flexibility. We demonstrate that the reserves offered are increased substantially if the inertia of a subset of the connected buildings is exploited too. Furthermore, we validate the control approach in a first-of-its-kind experiment on the actual system, where we show that while reserves are offered by the district system, the comfort in the connected buildings is maintained at all times.

We conclude the thesis in Chapter 6 and provide future research directions, which include scaling up data-driven control methods, data-driven modeling of multi-energy systems and physics-informed Machine Learning in the building domain.

### **1.3 Publications contributing to the thesis**

The work presented in this thesis mainly relies on published and forthcoming work, which has been created by me and my research collaborators. The specific journal and

conference research articles are listed below per chapter, sorted with ascending date of publication (or creation for working manuscripts).

## Chapter 2

- **Bünning, F.**, Bollinger, A., Heer, P., Smith, R.S. and Lygeros, J., 2019. Empirical validation of a data-driven heating demand simulation with error correction methods. in: Proceedings of Building Simulation 2019: 16th Conference of IBPSA, pp. 1428-1435.
- **Bünning, F.**, Heer, P., Smith, R.S. and Lygeros, J., 2019, November. Sensitivity analysis of data-driven building energy demand forecasts. Journal of Physics: Conference Series 1343, 012062.
- **Bünning, F.**, Heer, P., Smith, R.S. and Lygeros, J., 2020. Improved day ahead heating demand forecasting by online correction methods. Energy and Buildings 211, 109821.

## Chapter 3

- **Bünning, F.**, Warrington, J., Heer, P., Smith, R.S. and Lygeros, J., 2020. Frequency regulation with heat pumps using robust MPC with affine policies. IFAC-PapersOnLine 53, 13210-13215.
- **Bünning, F.**, Warrington, J., Heer, P., Smith, R.S. and Lygeros, J., 2020. Robust MPC with data-driven demand forecasting for frequency regulation with heat pumps. Submitted to Control Engineering Practice. arXiv:2009.06920.

## Chapter 4

- **Bünning, F.**, Huber, B., Heer, P., Aboudonia, A. and Lygeros, J., 2020. Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings. Energy and Buildings 211, 109792.
- **Bünning, F.**, Schalbetter, A., Aboudonia, A., Hudoba de Badyn, M., Heer, P. and Lygeros, J., 2021, May. Input convex neural networks for building MPC. in: Proceedings of the 3rd Conference on Learning for Dynamics and Control, PMLR. pp. 251-262.
- **Bünning, F.**, Huber, B., Schalbetter, A., Aboudonia, A., Hudoba de Badyn, M., Heer, P., Smith, R.S. and Lygeros, J., 2021. Physics-informed linear regression is competitive with two Machine Learning methods in residential building MPC. Accepted at Applied Energy.

## Chapter 5

- **Bünning, F.**, Heer, P., Smith, R.S. and Lygeros, J., 2021. Increasing electrical reserve provision in districts by exploiting energy flexibility of buildings with robust MPC. Working manuscript.

## 1.4 Other publications

The following papers have been prepared during my doctorate but are not part of this thesis.

- Saelens, D., De Jaeger, I., **Bünning, F.**, Mans, M., Vandermeulen, A., Van der Heijde, B., Garreau, E., Maccarini, A., Rønneseth, Ø., Sartori, I. and Helsen, L., 2019. Towards a DESTEST: a district energy simulation test developed in IBPSA Project 1. in: Proceedings of Building Simulation 2019: 16th Conference of IBPSA. pp. 3569-3577.
- **Bünning, F.**, Pfister, C., Aboudonia, A., Heer, P. and Lygeros, J., 2021. Comparing machine learning based methods to standard Regression methods for MPC on a virtual testbed. Presented at Building Simulation 2021: 17th Conference of IBPSA.
- Huber, B., **Bünning, F.**, Decoussemaeker, A., Heer, P., Aboudonia, A. and Lygeros, J., 2021. Benchmarking of data predictive control in a real-life apartment during heating season. Journal of Physics: Conference Series 2042, 012024.
- Derungs, C., Huber, B., **Bünning, F.**, Fricker, R., Stoller, S., Niesen, B., and Heer, P., 2021. Comprehensive energy demand and user behaviour data for building automation. Submitted to Nature Scientific Data.
- Lefebure, N., Khosravi, M., Hudoba de Badyn, M., **Bünning, F.**, Lygeros, J., Jones, C., Smith, R.S., 2021. Distributed model predictive control of buildings and energy hubs. Accepted at Energy and Buildings.

# Demand Forecasting with online-corrected ANN

Predictive building energy control on higher system levels, for example demand-side management in district energy systems, as will be presented in Chapter 3, often requires forecasts of the future energy demand of buildings or entire districts. Such forecasts can be done with data-driven methods such as Artificial Neural Networks. However, the prediction performance of Artificial Neural Networks suffers from high variance due to the stochastic nature of the training process. This means that two parameter-wise identical networks fitted to the same training data set perform differently well in forecasting the testing set. In this chapter, we propose two correction methods, one based on the forecasting error-autocorrelation, and one based on online learning, to obtain more reliable forecasting models.

The approach is tested in three case studies in the frame of day-ahead sub-hourly heating demand forecasting. In the first case study of a complex building, which has properties of a district heating system, it is demonstrated that the methods significantly reduce the variance in prediction performance and also increase the average prediction accuracy. When compared to grey-box and other black-box forecasting models, the approach performs better. In the second study, the approach is validated on three additional buildings where the performance of the approach is confirmed. In a third case study, the sensitivity with regards to quality and quantity of training data is investigated and it is demonstrated that the correction methods lower the amount of training data needed.

## 2.1 Introduction

In addition to energy management on a building level, there are a number of higher-level control tasks with the objective of optimizing energy consumption in the building domain or in coupled sectors (e.g. the electricity grid). Such tasks are for example the optimal control of the network temperature in low-temperature district heating and cool-

ing networks (Bünning et al., 2018), optimal control of energy hubs with various heating, cooling and storage technologies (Geidl et al., 2007; Arnold et al., 2009; Arnold and Andersson, 2011; Darivianakis et al., 2017), or frequency regulation with electrified heating systems and water storage tanks (Kondoh et al., 2011), as discussed in Chapters 3 and 5. Here, control tasks include minimizing energy costs, maximizing self-consumption, peak shaving and optimal management of long-term and short-term energy storage devices for example.

If predictive control is used to address these tasks, it is necessary to have detailed forecasts about the future energy consumption of individually connected buildings or the whole district. This has motivated an increasing research interest in demand forecasting for buildings for heating, cooling and overall electricity consumption, see for example the survey of different methods and models under (Mat Daut et al., 2017; Zhao and Magoulès, 2012; Harish and Kumar, 2016; Wang and Srinivasan, 2017; Fouquier et al., 2013; Suganthi and Samuel, 2012; Amasyali and El-Gohary, 2018; Ahmad et al., 2018). According to Amasyali and El-Gohary (2018), 31% of the related studies make forecasts with a yearly, monthly or daily resolution, while 57% consider hourly and 12% sub-hourly forecasts respectively, which would be suitable for the above named control tasks. Artificial Neural Networks (ANN) have been proven to perform well in a variety of complex and difficult tasks, such as gaming (Silver et al., 2017), computer vision (Wojna et al., 2018) and robotics (Tobin et al., 2017), and several studies have also demonstrated that ANN perform well on the task of short-term heating demand forecasting with high temporal resolution in the building and district domain.

On a district level, Kato et al. (2008) use a Recurrent Neural Network to make 24 hour ahead predictions of the hourly heating load of a real district heating system and compare it to a three-layered feed-forward ANN. The Recurrent Neural Network has an overall better forecasting performance, especially in periods that the authors consider as non-stationary. Similarly, Park et al. (2010) compare feed-forward ANN, Support Vector Machines (SVM) and the Particle Least Squares (PLS) method to perform 24 hour ahead predictions of the heating load of a real district heating system. PLS slightly outperforms SVM and ANN. Johansson et al. (2017) compare Extra-Trees Regressors (ETR) with feed-forward ANN on a real district heating testing case in Sweden for 24 hour ahead hourly forecast. ANN outperform ETR in this study. In Suryanarayana et al. (2018) the same Neural Network approach is compared to linear regression, ridge regression and lasso regression with an automatized feature selection process on two different district heating test cases in Sweden. ANN outperform the other three approaches in both cases. Saloux and Candanedo (2018) use ANN, SVM, and Regression trees (RT) and compare them to Piecewise Linear Regression on six hour ahead and 48 hour ahead heating demand forecasts with a sampling time of ten minutes in the case of a real district heating system. While ANN, SVM and RT clearly outperform the Piecewise Linear Regression, no distinguished performance difference is found between the other

methods.

On the individual building level, Paudel et al. (2014) use ANN together with a pseudo-dynamical model to make four-day predictions with a sampling time of 15 minutes for the heating demand of a French office building. Kwok and Lee (2011) and Leung et al. (2012) combine ANN with different occupancy prediction models to enhance the hourly load prediction of a large office building and a university building in Hong Kong respectively. The same forecasting task is performed for a commercial building in (Mestekemper et al., 2013). Here, the authors improve the prediction accuracy of ANN by splitting the cooling demand in a trend and a periodic signal and by training the ANN with the global optimization method Modal Trimming.

While all of the above named studies have promising results for energy demand forecasting with ANN, none of them communicate whether the achieved prediction results are average results or best case results, or if they were obtained with a fixed random seed. This is highly problematic as certain Machine Learning approaches and especially ANN suffer from high variance (Henderson et al., 2018; Recht, 2018; Bengio, 2012; Jamieson and Talwalkar, 2016): Depending on the random seed that is used while training the ANN, the ANN will have different prediction accuracies for the same training and testing data and same model parameters (i.e. number of neurons, number of layers, etc.).<sup>1</sup> This is due to the initialization of the network’s weights with random values and the following non-convex optimization that generally converges to a local minimum in the fitting process. Jovanović et al. (2015), Jetcheva et al. (2014) and De Felice and Yao (2011) use ensemble methods to overcome this problem in the context of building energy demand prediction. However, using ensembles also has disadvantages. For example, the prediction accuracy of an ensemble is not guaranteed to be better than the one of its best predictor. Moreover, using ensembles is computationally more expensive than using single predictors, which can obstruct their use in online control schemes. The computational effort typically grows linear with the number of predictors in the ensemble.

### 2.1.1 Contribution

We therefore make the following contribution with this chapter: A prediction model based on single ANN to make 24 hour ahead forecasts with a 15-minute sampling rate<sup>2</sup> of the heating demand for complex buildings and districts is developed. Two simple forecast correction methods, based on the error-autocorrelation and online learning, which make use of the history of previous forecasting errors during the online phase, are introduced.

The methods are verified in three case studies to examine the performance of the

---

<sup>1</sup>While fixing the random seed appears tempting, this would raise the question of which random seed to pick to reach the best performance.

<sup>2</sup>This is a common sample rate for many building energy management related control tasks.

resulting prediction methods, to test if they generalize to other buildings, and to evaluate requirements on the training data. This process is done to ensure that the methods are accurate enough to be used as disturbance forecasts for predictive controllers, to ensure that the methods are reusable for other buildings, as buildings are usually individual, and to verify that the required amount and quality of training data is suitable for practical application.

The first case study is based on data from the NEST building (Richner et al., 2017), which comprises independent modular control zones that allow it to mimic a district heating network. Here, we demonstrate that the methods reduce the ANN dependence on random initialization parameters and thus significantly reduce the variance in prediction performance. They also increase the average prediction accuracy and remove a prediction bias. Moreover, the methods are compared to other Machine Learning (or regression-based) approaches and to conventional resistor-capacitor (RC) building models, and are found to be competitive.

In a second case study, the generalization capability of the method is investigated by validating the approach on three additional buildings of different types and by comparing it to fitted 5R3C resistor-capacitor models. It is shown that the approach outperforms the RC model in all presented cases. Furthermore, improvements on the average accuracy and variance of the prediction accuracy of different instances of ANN are also observed for the other building test cases when the correction methods are applied.

In the third case study, we investigate the influence of sampling time, noise level and amount of available training measurement data as well as the quality of the weather forecast on the prediction accuracy. We demonstrate that the sampling time has a stronger influence on the prediction performance than the noise level and the amount of available data. Furthermore, we show that using measured ambient temperatures for training appears to provide no benefit compared to using weather forecasts.

## **2.1.2 Structure**

The remainder of the chapter is structured as follows: In Section 2.2, the forecasting task is introduced. In Section 2.3, our ANN based methodology and the forecast corrections, which are based on the forecasting error-autocorrelation and online learning, are described. In Section 2.4, the first case study is defined. The energy system under consideration is introduced, and the structure of the applied ANN and the benchmarking grey-box and black-box methods are defined. Then, the results regarding prediction accuracy and variance are discussed, as well as limitations of the study. Section 2.5 describes the configuration of the second case study first and then discusses the obtained results. Section 2.6 is structured analogously for Case Study 3. In Section 2.7, the chapter is concluded.



## 2.2 Problem statement

In this chapter the following forecasting task is considered: A heating demand forecast for a building or a district is made at midnight for the next 24 hours, sampled every 15 minutes, which is related to the reserve bidding mechanism used in Chapters 3 and 5, but also applicable for many other district energy management related tasks. The training and validation data are assumed to be sampled at the same rate.

## 2.3 Methodology

The forecast is made with a feed-forward ANN, which will be described in detail in Section 2.4.2. To model the closed-loop response of the building and its lower-level climate controllers to the ambient conditions, model inputs related to ambient conditions and time features are used. To reduce variance and to improve the forecasting accuracy, we introduce two methods that are used in the online phase of the forecasting task.

### 2.3.1 Error-autocorrelation correction

Empirical evidence suggests that forecast errors persist over a longer period of time than a single forecasting interval in the setting of building energy or district energy forecasting. For example, if windows are left open, a room temperature set point is changed or a fluid pump fails, this does not only have an effect on one single forecasting interval (15 minutes) but also on the following intervals, as the source of the error is usually not eliminated within one single interval.<sup>3</sup>

Motivated by this, we predict the error  $\tilde{e}$  for day  $D$  and forecast interval  $t \in [1, 96]$  by setting

$$\tilde{e}_{D,t} = \frac{1}{4} \sum_{\tau=93}^{96} e_{D-1,\tau} \Gamma_{ee}(t, E) \quad (2.1)$$

with

$$\Gamma_{ee}(l, E) = \frac{\mathbb{E}[(E - \mu)(E_{+l} - \mu)]}{\sigma^2}. \quad (2.2)$$

Here the first term in the right hand side of (2.1) denotes the average of the past day's forecasting errors  $e$  for the last four forecasting intervals; averaging over a few

---

<sup>3</sup>This is further demonstrated in section 2.4.4.

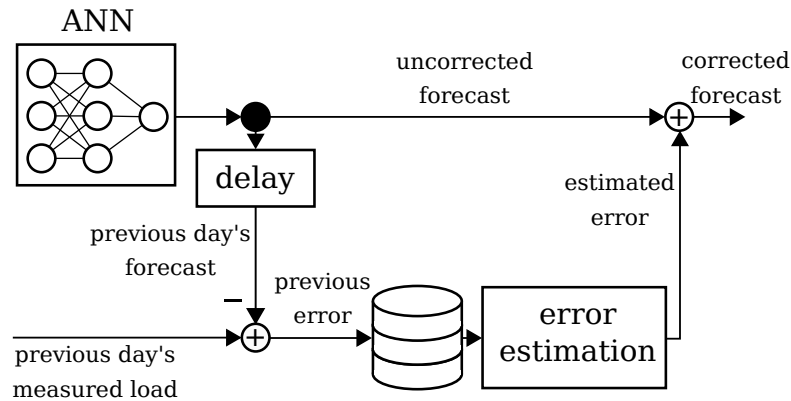


Figure 2.1: Scheme of forecast correction based on autocorrelation of forecasting errors

elements helps to denoise the error at the cost of a small decrease in the autocorrelation. The choice of four elements could be questioned, however, by further increasing the averaging interval, prediction power is lost because of the decreasing correlation of the error.  $\Gamma_{ee}(l, E)$  denotes the autocorrelation for lag  $l$ , which is calculated with the set  $E$  which contains all past forecasting errors  $e$  that have occurred until day  $D$  - including the errors of the training set.  $E_{+l}$  is the same set shifted by lag  $l$ . The expected value  $\mathbb{E}$ , the mean  $\mu$  and the standard deviation  $\sigma$  are empirical approximations for the stochastic process based on the set  $E$ .

Instead of including all past forecasting errors in  $E$ , one can also implement a shifting window and compute the autocorrelation using the forecast errors for the past few days. However, our experiments suggest that this does not improve the error prediction.

Figure 2.1 shows a schematic of the forecast correction procedure; A demand forecast is made with an ANN (uncorrected forecast). Based on the database of all previous forecasting errors and the measured error of the previous day's forecast, the error estimation described in (2.1)-(2.2) is made. The estimated error is then added to the uncorrected forecast to give a corrected demand forecast. At the end of the day, the difference between demand realization and demand forecast is added to the database. At the first prediction day in the online phase, the correction is based on the errors in the training set. The corresponding algorithm is shown in Algorithm 1.

### 2.3.2 Online learning

The second forecast correction method is online learning. Instead of using only the training set for training the model, after each day, the newly collected data (network inputs and demand realization) are used to retrain the ANN. The inputs are identical to the ones used in offline training. The argument behind this approach is again that phenomena in the building domain persist for time constants longer than hours and that the assumption "tomorrow will be like today" (also referred to as "persistence model" or

---

**Algorithm 1**

---

```

1: procedure ERROR-AUTOCORRELATION CORRECTION
2:   for  $D$  in runtime do
3:     for  $t$  in  $[1, 96]$  do
4:       forecast[D,t]:=ANN.predict(inputs[D,t])
5:       error[D,t]:=AvError[D-1]* $\Gamma_{ee}(t,E)$ 
6:       correctedFcast[D,t] := forecast[D,t]
           + error[D,t]
7:       wait until day passes
8:       measuredError[D] := realization[D]
           - forecast[D]
9:       E:=E.add(measuredError[D])

```

---

“naïve forecast”) has some validity here. By retraining the network on the previous day, we intentionally bias the network towards a certain solution, which is usually undesirable in Machine Learning (Huang et al., 2006), but potentially beneficial when the model is persistent. We can adjust this biasing via the learning rate of the optimizer used for training. This balances the importance between the training set and the daily feedback from the realised prediction errors. Combining this online training heuristic with the autocorrelation based error correction outlined above, leads to Algorithm 2.

It should be noted, that in line 10 a second forecast is made at the end of the day, after the ANN has been retrained on the inputs and demand realization of the day. The error of this second forecast is then added to the error database. This is done because next day’s error estimation needs to be performed for the retrained ANN. If this is not done, the error for the next day is overestimated.

---

**Algorithm 2**

---

```

1: procedure COMBINED CORRECTION
2:   for  $D$  in runtime do
3:     for  $t$  in  $[1, 96]$  do
4:       forecast[D,t]:=ANN.predict(inputs[D,t])
5:       error[D,t]:=AvError[D-1]* $\Gamma_{ee}(t,E)$ 
6:       correctedFcast[D,t] := forecast[D,t]
           + error[D,t]
7:       wait until day passes
8:       ANN:=ANN.train(inputs[D],realization[D])
9:       for  $t$  in  $[1, 96]$  do
10:        forecast2[D,t]:=ANN.predict(inputs[D,t])
11:        measuredError[D] := realization[D]
           - forecast2[D]
12:        E:=E.add(measuredError[D])

```

---



Figure 2.2: NEST building at Empa in Switzerland, Copyright: Zooley Braun - Stuttgart

## 2.4 Case study 1: General Validation

In this case study of a modern multi-use building in Switzerland, we demonstrate the general ability of the presented correction methods to increase accuracy and decrease the prediction variance between ANN with different initialization values. Besides discussing the individual effect of each correction method and analyzing the prediction error, we also compare the method to other black-box and grey-box demand prediction approaches. Here, we find that the presented approach outperforms all other tested methods.

### 2.4.1 System description

The NEST building at Empa in Switzerland (Richner et al., 2017), shown in Figure 2.2, is used as a case study. The building contains multiple individual units with different uses that can be added and removed from the building, in addition to office and meeting rooms that are permanently installed. At different times during the period covered by our data, two, three or four of these individual units were simultaneously in operation. One of the units was added during the training set (August 2017) and one during the test set (February 2018) of the case study. The configuration resembles a district heating system because different units have different use patterns, some of which resemble those of residential buildings while others those of office spaces. Moreover, all of the individual units are equipped with local controllers, whose details are unknown to the forecasting system. The demand forecast is therefore made for a collection of individual closed-loop energy systems. The units are connected via substations to a central heating system with a supply temperature of 38 °C and a return temperature of 28 °C. The heating system is served by a central heat pump with a maximum thermal capacity of 100 kW. A simplified schematic of the system can be seen in Figure 2.3. The heating demand that our algorithm is designed to forecast is the energy flow difference between supply

line and the return line. For more information on the complete system see (Lydon et al., 2017).

The building is set-up as a test bed for district heating and cooling and building energy experiments. It integrates approximately 540 sensors, whose data is stored in an SQL database with a temporal granularity of one minute, and approximately 200 controllable actuators. The measurement data used in this study contain 411 days (01.04.2017 to 16.05.2018) and include the measurements for ambient temperature and total heat consumption. The measured heating demand is shown in Figure 2.4. Figure 2.5 shows the ambient air temperature measured at the roof top of the building. Though, strictly speaking, weather forecasts should be used in demand forecasting, in this case study we use the actual ambient air temperature directly as an input. This is done to avoid the additional uncertainty in the inputs and to allow direct comparison of the different forecasting methods.

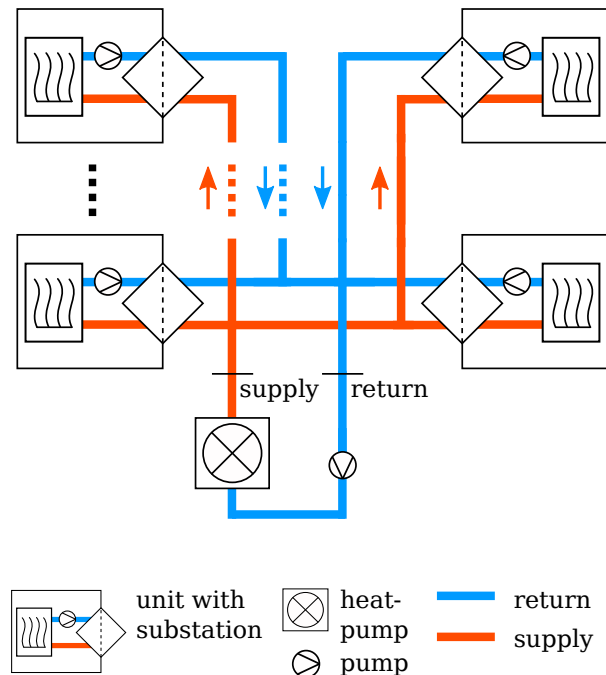


Figure 2.3: NEST heating system schematic. Orange lines denote the supply pipes and blue lines denote the return pipes. The individual units are connected to the heating system via heat exchangers.

## 2.4.2 Artificial Neural Network Configuration

We use feed-forward ANN with the Python package Keras (Chollet, 2018), which is a higher-level API to TensorFlow (Abadi et al., 2016). We refrain from fundamental introduction to the concept here and refer the reader to other sources (for example Shanmuganathan (2016); Basheer and Hajmeer (2000)). Other topologies, such as Recurrent




Figure 2.4: Measured total heating load of the NEST building. Periods where the trajectories are flat denote system outages, where the signal is held automatically by the database.




Figure 2.5: Measured ambient temperature at the roof of NEST building. Periods where the trajectories are flat denote system outages, where the signal is held automatically by the database.

Neural Networks (RNN) and Long-Short-Term Memory Networks (LSTMN) were investigated in preliminary studies but showed weaker performance. To fit the ANN to the training data, we use the optimizer Adam (Kingma and Ba, 2017) with standard parameters. This includes the learning rate of 0.001 for both offline and online learning. As a loss function, the mean squared error and as activation functions, Rectified Linear Units (ReLU) (Nair and Hinton, 2010) are used.

## Architecture

As we want to forecast the heating demand in a day-ahead fashion, we forecast 96 data points at once. (One value every 15 minutes for one day.) For the structure of the ANN this gives two possibilities: 1) The network is built with 96 outputs and the forecast for the day is performed in one single step. 2) The network is built with one output and the forecast is performed with 96 different input vectors to forecast 96 single output values.

Option 1 has one major advantage: As the heating demand (forecast target) is highly correlated for small time delays, the accuracy of the prediction is very precise for the first few forecast data points at the beginning of the day. This is because the network

makes use of the measured heating demands at the end of the previous day. In option 2, this is not possible as, for example, we cannot use the demand of  $t - 4$  as an input for the network because for forecasts  $t = 5, \dots, 96$  this input has not yet been measured at the time the forecast is made. As an alternative, the forecasts of earlier time steps could be used as inputs, but this would potentially lead to error propagation. However, in case of option 1 the network needs a much higher complexity in terms of number of layers and nodes, and thus potentially a larger training set compared to option 2. Preliminary forecast runs have shown that this results in reduced prediction performance. For option 2, the disadvantage of not being able to use the last measured heating demands of the previous day as inputs is compensated by the error-autocorrelation based correction method that was introduced in Section 2.3.1. We therefore use option 2 in this study.

### Feature selection

Although NEST also has measurements for wind speed and solar radiation, the ambient temperature was selected as the only weather-related feature<sup>4</sup>, because preliminary experiments suggested that the other measurements do not substantially improve forecasting accuracy. This might appear surprising as generally solar radiation has indeed an effect on the heating demand of a building, however, with a Pearson correlation coefficient of 0.54 between ambient temperature and solar radiation in the considered data set, it appears that the ambient temperature incorporates much of the information of the solar radiation already.

With the assumption that the demand of the system shows daily and weekly patterns, as the units are used for living and as office spaces, two more features were added: the total heat consumption with a time delay of one day (again following the assumption "tomorrow will be like today") and the total heat consumption with a time delay of one week (motivated by the assumption that weekly routines can be expected). This allows the prediction of the energy demand today based on the demand at the corresponding time yesterday and last week. Finally, the time stamp of all features was implemented as two inputs: hour of the day and weekday/weekend<sup>5</sup>. Hour of the day was one-hot encoded (Aggarwal, 2018), meaning that instead of one continuous input 24 binary inputs are used. Preliminary experiments suggested that this improves the forecast accuracy. We assume that the occupancy related internal gains are indirectly captured by the time-related inputs and do not use a separate occupancy model. The resulting network has 28 inputs, of which three are continuous and 25 binary, and one continuous output.

---

<sup>4</sup>The term "feature" can be regarded as a synonym for "input", but is the more common term in Machine Learning.

<sup>5</sup>Weekday/weekend was chosen instead of working day/non-working day as the difference is marginal and the implementation effort substantially lower.

Table 2.1: ANN parameters

Parameter	Value
Number of hidden layers	2
Number of nodes per layer	8
Number of Epochs	10
Input scaling method	Linear [-1,1]

### Parameter selection

There is no proven rule on how to best determine the parameters of ANN. These values are commonly set based on heuristics and an iterative approach that involves a substantial amount of trial and error. Generally, a trade-off between overfitting and underfitting of the training data has to be found.

A preliminary sensitivity analysis was conducted to decide on the number of hidden layers, number of nodes per layer, number of training epochs (number of times the network's weights are updated based on each individual sample in the training set), and on the input scaling method. The results for the suggested optimal parameter values are given in Table 2.1.

### 2.4.3 Forecasting methods for benchmarking

To assess the performance of the ANN with the introduced correction methods for forecasting the thermal demand of NEST, their prediction accuracy is compared to those of other state of the art prediction methods. These are grey-box models in the form of resistor-capacitor building models, and black-box models in the form of other regression-based or Machine Learning based approaches.

#### Grey-box method (RC models)

For the grey-box approach, resistor-capacitor (RC) models are used. We note that while resistor-capacitor and other grey-box models are often used as optimization models for selecting control inputs for building level Model Predictive Control, they are also used for forecasting (De Coninck et al., 2015; Zhou et al., 2008). Moreover, grey-box and white-box models are also commonly used for building performance simulation, which is essentially forecasting without using the forecasts online, see for example (Zhao and Magoulès, 2012).

Four different topologies were considered as shown in Figure 2.6, inspired by the models used in (De Coninck et al., 2015). In this approach, the thermal capacity of wall materials, floor materials and air volumes is represented by capacitors, whereas the thermal resistance of walls and floors is represented by resistors. The parameters



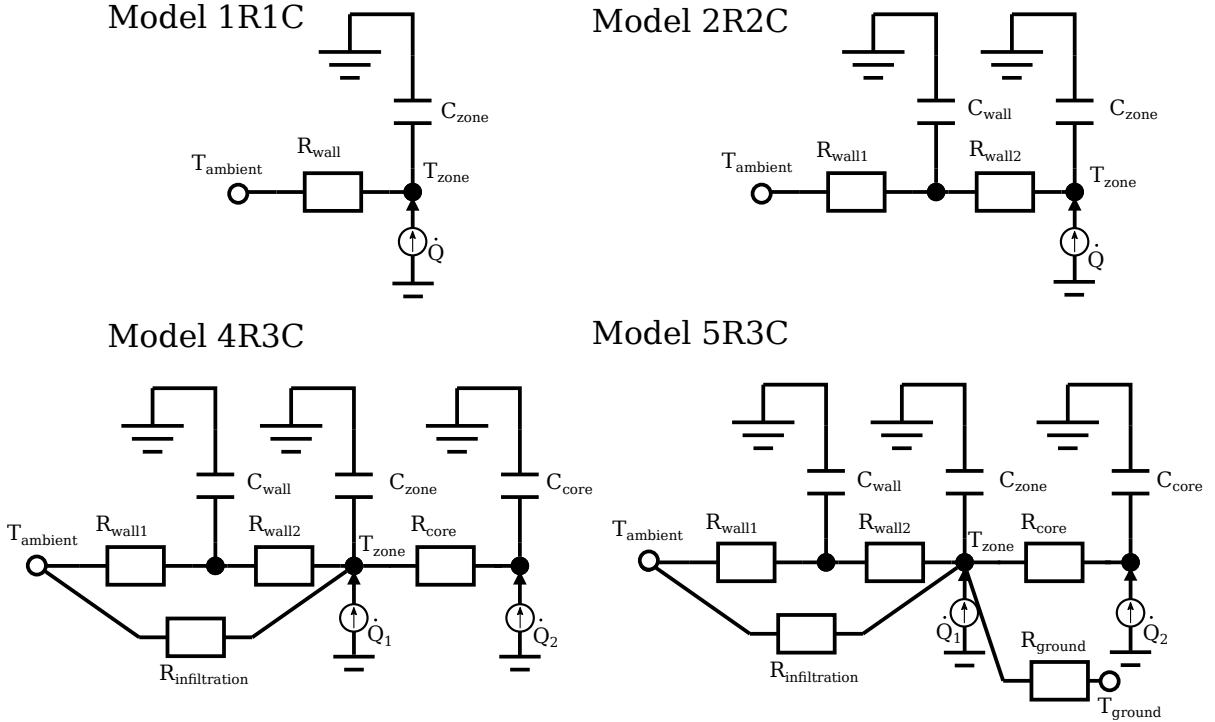


Figure 2.6: RC building models with varying numbers of resistors and capacitors.

of all resistors and capacitors are estimated to fit a heating demand measurement time series. As the RC models are purely thermal models, the ambient temperature is the only used input. The models were implemented in the modeling language Modelica (Mattsson et al., 1998) and simulated in Dymola (Brück et al., 2002). To estimate the parameters, a Covariance matrix adaptation evolution strategy (CMA-ES) (Hansen et al., 2003) optimizer was used in Python. CMA-ES is a evolutionary solver to find the global optimum of non-linear and non-convex optimization problems. It is considered the state-of-the-art algorithm in evolutionary computation (Loshchilov and Hutter, 2016).

In the 1R1C and 2R2C models, the parameters were estimated such that the mean squared error between the heating energy  $\dot{Q}$  applied to the model and the target heating demand is minimized. In model 1R1C,  $\dot{Q}$  was set equal to the heat flux through the wall resistor  $R_{\text{wall}}$  and in model 2R2C,  $\dot{Q}$  was set equal to the heat flux through wall 2,  $R_{\text{wall}2}$ , which keeps the temperature of  $C_{\text{zone}}$  at a constant 20 °C.

As an additional capacitor for the building core,  $C_{\text{core}}$ , was added for the models 4R3C and 5R3C, the mapping of the target heating demand is more complex. The optimization was set up to minimize the mean squared error between the target heating and  $\dot{Q}_{\text{total}}$ ,

$$\dot{Q}_{\text{total}} = \dot{Q}_1 + \dot{Q}_2, \quad (2.3)$$

which is the sum of the heat distributed via fast emission systems (e.g. air conditioning) to  $C_{\text{zone}}$  ( $\dot{Q}_1$ ) and the heat distributed via slow emission systems (e.g. floor heating) to  $C_{\text{core}}$  ( $\dot{Q}_2$ ). Moreover, as the direction of heat loss is not unique in these models (losses can go from  $C_{\text{zone}}$  to  $C_{\text{core}}$  and/or to  $T_{\text{ground}}$ ), a P-controller was set-up to keep the temperature of  $C_{\text{zone}}$  at a constant  $T_{\text{set}} = 20$  °C:

$$\dot{Q}_{\text{total}} = k_p \times (T_{\text{set}} - T_{\text{zone}}). \quad (2.4)$$

A P-controller is chosen because buildings are often controlled by thermostatic valves, which approximately have proportional controller behaviour. Moreover, the closed-loop demand response to the ambient temperature is assumed to be dominated by the building dynamics, which renders the choice for the assumed controller less important.<sup>6</sup> Additionally a fraction coefficient  $c_f$  was introduced to distribute the heat between air conditioning and floor heating, such that

$$\dot{Q}_1 = c_f \times \dot{Q}_{\text{total}} \quad (2.5)$$

and

$$\dot{Q}_2 = (1 - c_f) \times \dot{Q}_{\text{total}}. \quad (2.6)$$

The values of  $k_p$  and  $c_f$  are also determined in the fitting process.

### **Black-box methods (regression-based methods)**

To implement the benchmark black-box methods, we use Python's scikit-learn package (Pedregosa et al., 2011). The package features a variety of different regression or Machine Learning methods. In the case study, the following methods are used:

1. Least squares Linear Regression
2. Support Vector Machine
3. Huber Regressor
4. Orthogonal Matching Pursuit
5. SGD Regressor

---

<sup>6</sup>This is the case for any well-regulated building.

## 6. Decision Tree Regression

## 7. Random Forest

The data sets used for fitting and testing are identical to those used for the ANN method, including the scaling of inputs and outputs and the one-hot encoding of categorical features. For all methods, a grid search regarding model tuning parameters was performed to improve the forecasting results.<sup>7</sup>

### 2.4.4 Results and discussion

For the case study, the measurement data from NEST is divided into a training set and a testing set. The training set consists of 70% of the whole data set, which corresponds to 287 days (27552 data points). The testing set consists of 124 days (11904 data points).

The data was post-processed with the help of the pandas package (McKinney, 2011) in Python 3. Missing data points were first linearly interpolated or extrapolated from neighbouring data points. This excludes data points that have already been handled by the NEST database (e.g. the constant values around September 2017 in Figure 2.5). The data set was then re-sampled to a 15-minute interval using the mean value of all relevant data points.

To benchmark the prediction performance, the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{i \in N} (y_i - f(x_i))^2}{\sum_{i \in N} (y_i - \bar{y}_N)^2}, \quad (2.7)$$

is used. It becomes *zero* if the forecast  $f(x_i)$  is as good as taking the average  $\bar{y}_N$  of the data in the considered set  $N$  as a forecast, and becomes *one* if the forecast exactly matches the validation data  $y_i$ . In principle,  $R^2$  may also become negative if the forecast is worse than taking the average, though this was not observed with any of the methods tested here.

We train and forecast with 100 individual ANN as the initial input weights of each node in the network are initialized with a random value. This randomness leads to different prediction performance for each network as the training process involves the solution of a non-convex optimization problem that generally converges to a local minimum. By simulating 100 networks, a more confident statement about  $R^2$  can be made and the quartiles can be studied.

---

<sup>7</sup>These parameters are in particular *kernel*, *gamma*, *C* and *degree* for Support Vector Machine, *epsilon* and *alpha* for Huber Regressor, *nonzero coefficients* for Orthogonal Matching Pursuit, *alpha* and *L1 ratio* for SGD Regressor, *depth* for Decision Tree Regression, and *number of estimators* and *depth* for Random Forests. We refer the reader to the API reference of scikit-learn for details on the tuning parameters.

### Prediction accuracy, quartiles and improvements through forecast corrections

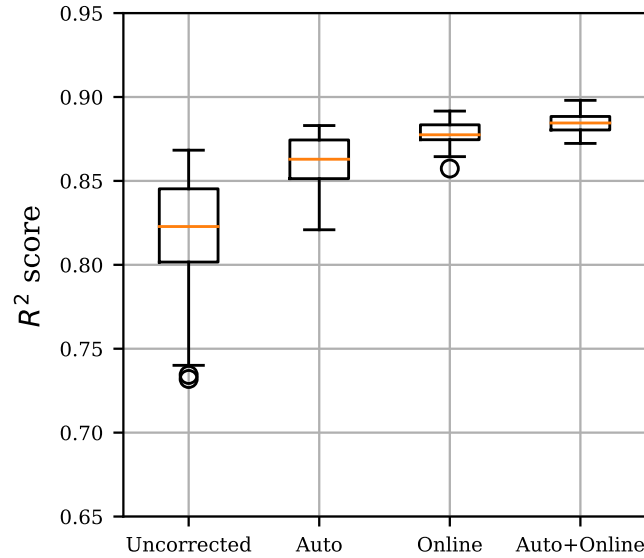


Figure 2.7: Coefficient of determination for ANN and forecast correction methods for 100 individual networks, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median.

Figure 2.7 shows the box plot of the coefficient of determination in the test set of 100 different networks for the uncorrected ANN, ANN with error-autocorrelation correction, ANN with online learning and ANN with error-autocorrelation + online learning for full day-ahead forecasts. First, it should be noted that the variance in performance for uncorrected ANN is indeed very high, as discussed in Section 2.1. The highest performing network reaches  $R^2 = 0.868$ , while the lowest performing network reaches  $R^2 = 0.732$ . It can be seen that the quartiles regarding prediction accuracy amongst different network instances are significantly reduced when correction methods are used. While the interquartile range (IQR) is 0.038 in the uncorrected case, it reduces to 0.023 in the case with error-autocorrelation correction, 0.009 in the case with online learning and 0.008 in the case where both correction methods are used. With both corrections in place, all  $R^2$  lie between 0.872 and 0.898. The average achieved  $R^2$  are 0.818, 0.860, 0.878 and 0.885 subsequently. Figures A.1 - A.3 in the Appendix show the results of a second experiment with 100 runs for the variance of three different KPI, Mean Squared Error (MSE), Mean Absolute Error (MAE) and Coefficient of Variation of the Root-Mean Squared Error (CV RMSE). The trend is identical to the one observed for  $R^2$  in the first experiment. For clarity, the analysis is therefore limited to  $R^2$  in the following.

Figure 2.8 shows the coefficient of determination for a 2 hour ahead prediction and

Figure 2.9 for a 24 hour ahead prediction<sup>8</sup>, again for 100 networks. First, it can be noticed that the performance in terms of median and IQR is similar for the uncorrected forecasts. However, in case of the corrected forecasts, the performance is better in the 2h ahead case, because the forecast is closer to the model update and to the last measured error, on which the error-autocorrelation correction is based on. Moreover, while the correction based on online learning shows a better performance in terms of median for the full forecast (Fig. 2.7) and 24h ahead forecast, the correction based on error-autocorrelation is more effective in the 2h ahead case. However, the variance is significantly reduced by both correction methods in all cases.

These results suggest that the corrections introduced in Section 2.3 can to a large extent alleviate a main disadvantage of ANN, namely the dependence of the prediction performance on randomly initialized node input weights, without the necessity of using an ensemble method. Moreover, the average prediction performance is significantly improved.

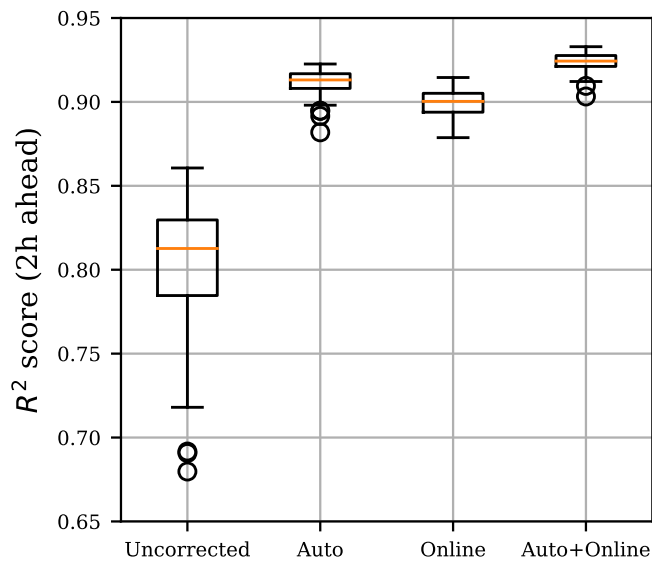


Figure 2.8: Coefficient of determination of 2h ahead prediction of 100 ANN and forecast correction methods, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median.

<sup>8</sup>In contrast to the full day-ahead forecasts, the 24h ahead prediction only includes the sample for exactly 24h ahead, but not the samples between the forecasting time and 24h. I.e. it does not include the samples for 15m ahead to 23h45m ahead.




Figure 2.9: Coefficient of determination of 24h ahead prediction of 100 ANN and forecast correction methods, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median.

### Influence of training set on prediction performance

The training set and the online application of forecasts can often fall into different seasons for real-life applications, which results in different probability distributions of samples in training and application (or in case of this study in the testing set). To investigate the influence of different distributions, we have trained 100 networks based on randomly sub-sampled sets, with random length and random location in the original set, from the original training set described above. The networks are validated on the same testing set for comparable  $R^2$ .

Figure 2.10 shows the variance of the coefficient of determination. It can be seen that both the median and IQR have significantly worsened for the uncorrected networks compared to Figure 2.7. In the case where both correction methods are applied, both median (0.868) and IQR (0.0183) are close to the ones achieved with the full training set.

### Prediction error analysis

To further evaluate the impact of the correction terms introduced in Section 2.3 we analysed the statistics of the forecast error with and without the corrections.

Figure 2.11 shows the autocorrelation of the error for one instance of an uncorrected




Figure 2.10: Coefficient of determination of 100 ANN and forecast correction methods with randomized training period, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median.

ANN forecast (with a relatively high  $R^2$  of 0.841) and the remaining error after correction for one instance of a corrected ANN with error-autocorrelation and online learning (with  $R^2$  of 0.887). It can be seen that the signal for the corrected ANN is less correlated than the uncorrected signal. However, the residual is still not “white”, as the autocorrelation is still substantial for many values of the lag. We conjecture that this residual autocorrelation is an intrinsic artefact of the day-ahead forecast: Disturbances that occur during the forecasting period will introduce a systematic forecasting error for the rest of the period, as there is no chance to correct the forecast until the next forecasting period begins.<sup>9</sup> Moreover, one can see that there is no peak at  $t = 96$ , i.e. after one day. This suggests that the network captured all regular daily time-dependencies of the demand and has no systematic error regarding this aspect.

Figure 2.12 compares the forecasting error histogram of one instance of ANN without correction to the one using both error correction methods. The plot is cut at a frequency of 10 for the sake of readability, because by far most errors are 0 for both cases. The interquartile range reduces from 0.054 in the uncorrected case to 0.033 in the corrected

<sup>9</sup>To use the example of an opened window again: If the window is opened at day  $D - 1$ , we can use the resulting forecast error in day  $D - 1$  in our forecast corrections for day  $D$ . However, if the window is opened at day  $D$ , for example at  $t = 5$ , this will introduce a systematic error in the forecast for  $t = [6 \dots 96]$ . This error will show up in the autocorrelation plot and is impossible to avoid in this forecasting setting.

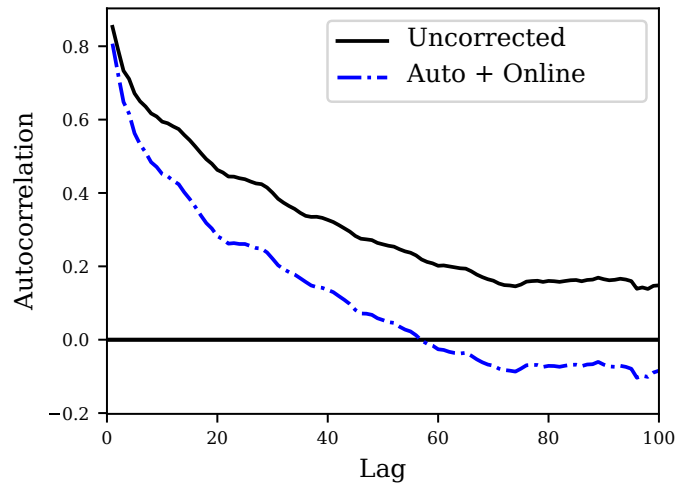


Figure 2.11: Autocorrelation of the residual uncorrected ANN (black), and ANN with correction by error-autocorrelation and online learning (blue).

case. Furthermore, the median improves from -0.025 to 0, which shows that the bias in the forecast is effectively eliminated by the forecast corrections.

### Forecasting trajectory examples

Figure 2.13 shows examples of the forecast heating demand trajectory at different temporal resolutions for an ANN that includes the error-autocorrelation and online learning corrections. The figure also shows the empirical confidence bounds for the forecast computed with a posteriori analysis of the absolute error distribution of the testing set. The dark-grey background depicts a confidence bound of 68% (one standard deviation) and the light-grey background depicts a confidence bound of 95% (two standard deviations). Such confidence intervals could be useful if one uses the forecast for predictive building energy management based on robust or stochastic model predictive control, as done in Chapter 3. For certain control tasks, such as temperature control within a building, these confidence bounds might appear large. However, for high level energy management tasks, methods such as affine decision rules (see (Goulart et al., 2006; Warrington et al., 2012, 2014) and Chapter 3) can be deployed to introduce recourse in receding horizon control schemes that significantly reduce the impact of forecast uncertainty on the control performance.

The top graph of Figure 2.13 shows the prediction and the actual heating demand for the complete testing set. The general trend is captured well. We note in passing that the model seems to cope well with another unit being added to the building in February (as described in Section 2.4.1), as the general trend is matched equally well before and after February. The middle graph shows a more detailed view of five days with a medium heating demand in March. At this scale, the capabilities and limitations



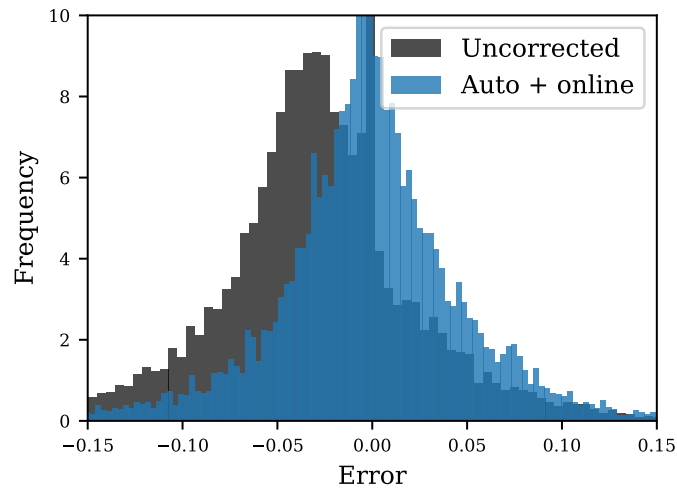


Figure 2.12: Forecasting error histogram for uncorrected ANN (black), and ANN with correction by error-autocorrelation and online learning (blue). The y-axis is limited to 10 for improved visibility. The number of bins is set according to the Freedman-Diaconis rule.

of the forecast correction methods can be observed. At the beginning of the 18th of March and 19th of March, the forecast signal is corrected towards the true demand, because an error was observed at the end of the previous days. The time of correction is indicated by the vertical grey lines. On the other hand, in the middle of 17th of March the forecast deviates from the actual demand for several hours. However, this error can not be corrected, as the correction is applied only at the beginning of each 24 hour forecasting interval. In the bottom graph, a system shut-down or failure occurs in the afternoon of the 26th of February. The forecast is not corrected until the start of the next day. At the beginning of the 28th of February and 1st of March a correction of the forecast towards the real demand at the beginning of the day can be observed again.

### Comparison of different forecasting methods

Table 2.2 compares the coefficient of determination in the testing set of the ANN with both correction methods applied and the benchmark forecasting methods for full 24 hour forecasts.<sup>10</sup> It can be seen that the ANN, corrected with both correction methods, outperform the other forecasting methods.

Decision Tree Regression and Random Forest reach a high  $R^2$  without any additional corrections. It is possible that with appropriate correction methods (similar to the ones introduced in this study) the performance of our corrected ANN can be reached or exceeded. With respect to the correction methods introduced here, we note that online

<sup>10</sup>N-step ahead error metrics are not analysed because the models are not updated or corrected and the inputs are real measurements and not forecasts.




Figure 2.13: Forecasting examples (orange) compared to the real heating load (blue) with varying temporal resolution. The time of correction is indicated by the vertical grey lines. The grey shadings indicate confidence bounds of 68% and 95%.

Table 2.2: Coefficient of determination for all tested heating demand forecasting methods.

Method	$R^2$
ANN online + autocorr.	0.885
Huber Regressor	0.768
Orthogonal Matching Pursuit	0.770
Least squares Linear Regression	0.770
SGD Regressor	0.775
Support Vector Machine	0.811
Decision Tree Regression	0.844
Random Forest	0.860
Model 1R1C	0.691
Model 2R2C	0.702
Model 4R3C	0.739
Model 5R3C	0.761

learning is not a viable option for regression trees and random forests, because of the way these are constructed. Both methods are based on decision trees in which the input variables are split in partitions and the output is approximated with a constant value in these partitions (see Chapter 4). The splitting variables and splitting points are chosen heuristically. In online learning with ANN, the network weights are updated with the measured input-output samples every day. The importance of the new samples can be adjusted with a learning rate. Such an adjustment of splitting variable and point is not possible in trees, as the whole tree would have to be rebuilt from the changed splitting point down. Thus, regression trees or random forests always have to be trained on the complete data set and cannot be updated with single samples. Retraining on a combined set of historical data and daily measurements is in theory possible, but it would not bias the forecast towards the recent behaviour of the building. Corrections based on error-autocorrelation could, however, be used to correct forecasts made by any forecasting method.

The resistor-capacitor based models perform worse than the majority of the regression based methods in this case study. This could be due to the fact that these models can only make use of the ambient temperature as an input, making it impossible to capture time-dependent phenomena, such as occupancy for example. This shortcoming could be improved by an occupancy model. Moreover, these models involve only a small number of parameters, in the form of resistor and capacitor values. Increasing the model complexity further could improve forecasting, as suggested by the fact that the  $R^2$  of the RC models in our results continue to increase as more parameters are added. The addition of more RC components of course implies increased modeling effort. Finally, RC models are often used in this context for modeling building thermal dynamics (Sturzenegger et al., 2014) and not for demand forecasting. This study suggest that they are not equally suitable

for the latter.

### Limitations

As evident from the discussion in the previous section, the case study has certain limitations. The scope of the study is to find correction methods that solve the problem of prediction accuracy variance of ANN and to increase their prediction performance. As KPIs of prediction performance parameters are not comparable between different studies in literature and in this work (as they heavily depend on the predicted trajectory), ANN combined with online correction methods were benchmarked against other forecasting methods. As the scope of this study is not a comprehensive review and comparison of forecasting methods, the benchmarking methods have certain limitations and could potentially be improved in further studies. For example, besides the already mentioned measures, the prediction performance of the grey-box models could benefit from re-initializing the model states with temperature measurements from the building at the beginning of each forecasting period. However, there are also arguments that speak against this. In particular, temperature measurements of individual buildings are not always available in an energy-hub or a different higher level control context, for example because of households having privacy concerns regarding sharing such measurements with the system operator, or simply because of the lack of installed sensors.

## 2.5 Case study 2: Transferrability to other Buildings

The previous case study has demonstrated that ANN with online correction methods show high prediction accuracy on a modern multi-use building that resembles a small district heating system, and outperform other regression-based methods such as random forests and support vector machines, as well as fitted resistor-capacitor models with varying complexity. Moreover, the dependence of ANN prediction performance on initialization parameters, which are commonly randomly set, is significantly lowered with the shown methods.

However, with respect to the viability of the methods for practical application, it should be demonstrated that both the concept, and the found ANN architecture and input features can be applied to more than one specific building. In the second case study, we therefore aim to increase confidence in the methods by applying them to three additional buildings and by comparing the prediction accuracy to a fitted 5R3C building model.

It is shown that the approach outperforms the RC models in all presented cases. Furthermore, it can be seen that an improvement on the prediction accuracy and a reduction of the variance of the prediction performance of different instances of ANN

holds for all test cases.

### 2.5.1 General Description

The same forecasting task as in the previous case study is assumed: A heat demand forecast is made at midnight for the next 24 hours, sampled every 15 minutes. The training and validation data are assumed to be sampled at the same frequency. Furthermore, perfect knowledge of the ambient temperature for the forecasting period is assumed.

To validate the transferability of the ANN with online correction methods to other buildings, additionally to the NEST building, three other buildings were selected for the case study. For each building, a dataset of 15-minutely heat demand measurements and ambient temperature measurements were available with lengths of one to three years. In each case, the first 70% of the available dataset were used for model training (for both ANN and RC model) and 30% of the dataset were used for model validation. The buildings are introduced in the following and are depicted in Fig. 2.14.

**Building (a)** is the NEST building, as described in Case Study 1 (Section 2.4). 13.5 months of measured demand data were used for this building.

**Building (b)** is the ETL building located at the zentrum campus of ETH Zürich. It is an eleven-story office building that harbours several research institutes and also experimental laboratories. Six stories lie below ground. 36 months of measured demand data were used for this building.

**Building (c)** is the Bauhalle at Empa. The building has offices and testing facilities for material science experiments. The offices are spread on three stories while some of the laboratories have the ceiling height of these three stories. The building also has a basement. 24 months of measured demand data were used for this building.

**Building (d)** is the Verwaltungsgebäude at Empa. It is an office building with some presentation rooms. It features three stories. 24 months of measured demand data were used for this building.

Both model types, ANN with online correction methods, and the 5R3C resistor-capacitor models were configured and parametrized in the same way as described in Case study 1 (Section 2.4).

### 2.5.2 Results and discussion

#### Variance

Figure 2.15 (a)-(d) shows the variance of the prediction performance in terms of  $R^2$  for 100 different instances of neural networks for each building. The first column depicts uncorrected networks, the second column networks with error-autocorrelation correction,



(a) NEST at Empa, © Zoey Braun.



(b) ETL at ETH Zurich.



(c) Bauhalle at Empa, © Heinrich Helfenstein.



(d) Verwaltungsgebäude at Empa.

Figure 2.14: Buildings used in the validation study.

the third column networks with online learning and the last column a combination of both correction measures.

Subplot (a) shows the results from the same building that was used in Case Study 1: the median of the coefficient of determination increases from the case where no corrections are made to the case where both correction methods are combined. Moreover, also the variance reduces. The figure shows that the same trend occurs in the cases of buildings (b) and (c). Online learning gives the bigger improvement when compared to the correction based on the error-autocorrelation. However, the combination of both gives the best results in those cases. While the variance does not improve much, the median of the coefficient of determination further increases with the combined method.

In the case of building (d), the trend is not as strong but still present. The median increases with each individual correction method, but the median  $R^2$  of the online learning correction is lower than the one for error-autocorrelation correction. The median of the combined method is again the highest one. The maximum achieved coefficient of performance is lower for the online learning and the combination of both correction




Figure 2.15: Variance of the coefficient of determination for all tested buildings. Boxes describe the interquartile range, whiskers denote 1.5 times the interquartile range, circles indicate outliers and the orange line indicates the median.

methods compared to the uncorrected network and the correction method based on error-autocorrelation. This is due to overcorrection of already well predicting networks.

The difference in the result between building (d) and the other buildings could be explained with the quality of the used data set. The measured heating demand of building (d) has a quantization of 0.5 kW at a maximum load of 26.5 kW., while the “next worst” building’s measured demand has a quantization of 0.5 kW at a maximum load of 86 kW.

The results generally support the findings of the previous case study and indicate that the correction methods work for more than one single building. It could be questioned why the combination of both correction methods in most cases gives a further improvement over the individual methods as they both address a similar problem: the miss-match between forecast and realization due to recently changed behaviour of the

Table 2.3:  $R^2$  of different modeling approaches for all tested buildings.

Building	ANN	ANN auto	ANN online	ANN auto+online	5R3C model
Building (a)	0.818	0.860	0.878	0.885	0.761
Building (b)	0.886	0.907	0.933	0.936	0.890
Building (c)	0.674	0.752	0.793	0.809	0.676
Building (d)	0.856	0.858	0.860	0.862	0.747

building. However, the methods correct slightly different phenomena. The correction based on error-autocorrelation can react very efficiently to errors that have occurred shortly before the end of the last day’s forecast and apply efficient correction at the beginning of the next day’s forecast, because the confidence in the error estimation is high (because the autocorrelation is high). However, strong corrections are only possible for the first few intervals of the day, as the autocorrelation of the forecasting error decreases during the course of the day. The correction based on only learning corrects to a lesser extent on each individual interval, as the ANN still needs to be able to generalize. However, this correction can be applied over the course of the full day’s forecast and not just at the beginning of the day. The combination of both therefore gives the best result.

## Accuracy

Table 2.3 shows the mean coefficient of determination (as opposed to the median discussed before) of the different ANN approaches for all buildings as well as the coefficient of determination achieved by the 5R3C building model, which is fitted as described in Section 2.4.3. It can be seen that the average  $R^2$  increases for all four buildings when the individual correction methods are applied. A combination of both approaches gives a further improvement in all cases. This is consistent with the results of the previous case study. Moreover, the fully corrected ANN outperform the RC models in all tested cases.

The results do not indicate that the approach outperforms RC building models in every possible building simulation case. The RC model could for example be made more complex or the parameter fitting could be improved, both of which methods could lead to higher coefficients of determination. However, the results give an indication that the ANN perform reasonably well compared to conventional methods and generalize to other buildings.<sup>11</sup>

<sup>11</sup>Note that, between buildings it is not possible to compare the coefficient of determination because it depends on the distribution of heating loads in the data set, which is different for each building.



### Physical behaviour of the forecast

Figure 2.16 shows the forecast and real heating load of building (b) from an excerpt of the testing set. First, it can be seen that the real load is captured reasonably well by the forecast, although the peaks in the last third of the excerpt are not predicted well. However, to assess the quality of the forecast, the coefficient of determination is more descriptive.




Figure 2.16: Forecasting trajectory example from building (b) test set. The orange line denotes the forecast and the blue line denotes the true heating demand.

Secondly, the plot shows that the forecast does not show any non-physical behaviour in the sense of unreasonably high or low forecasting values. The results for all other buildings allow the same observation. (The full forecasting plots of both training and testing set for one instance of an ANN with both correction methods applied to all four buildings is shown in the Appendix A.2)

This observation is made with respect to the fact that the extrapolation capability of ANN is sometimes considered weak. Although these exemplary results cannot guarantee that non-physical behaviour is impossible with other forecasting inputs or for other building cases, the positive results for all four buildings indicate a robustness that is sufficient for non-safety-critical control tasks. To further increase the confidence in the forecast, the signal could be limited in its derivative and maximum/minimum values.

## 2.6 Case study 3: Sensitivity with respect to data quality and quantity

In the first case study (Section 2.4), we have demonstrated that online-corrected ANN have superior prediction accuracy compared to other Machine Learning methods and to a variety of fitted simple resistor-capacitor models. Furthermore, the correction methods allowed a significant reduction of variance in the prediction performance of ANN. In the second case study (Section 2.5), the method was validated on four individual buildings with different uses, availability and quality of measurement data. The correc-

tion methods significantly improved the forecast quality and reliability and the approach outperformed a fitted 5R3C resistor-capacitor building model in all cases. However, the results indicated that the prediction performance could be dependent on the quality of the training data.

In a third case study, we therefore investigate the dependence of the forecasting approach on the sampling time and the noise of the measurement data, as well as the available amount of data and the quality of the weather forecast used as an input. The results indicate that the sampling time of measurements has a larger impact on the prediction quality than noise levels, and that one week of historical training data is sufficient for meaningful demand predictions if online correction methods are applied. Moreover, using ambient temperature measurements for ANN training appears to offer no benefit compared to using weather forecasts.

### 2.6.1 General description

The same forecasting task is assumed as in the previous case studies of this chapter. Moreover, the NEST building at Empa in Switzerland with the same type of input data as described in Case Study 1 is used as the building to be investigated. The ANN are also configured as described in Case Study 1. Measurement data of the whole building's heating demand over a period of 13.5 months is used. To investigate the sensitivity with regards to data quantity and quality, the input and output data is manipulated as follows.

**Data preparation for varying sampling time:** We assume that the true measurement data is sampled in 15-minute intervals. To simulate different sampling times, we sub-sample the data set to 30 minutes, 60 minutes and 120 minutes by taking the mean over all relevant samples. The window to compute the mean starts with the sample under consideration, for example: in the 30-minute sub-sampled set, the data point for 2017-04-01 00:00 is calculated by taking the mean of samples 2017-04-01 00:00 and 2017-04-01 00:15 from the 15-minute data set.

**Data preparation for varying measurement noise:** We assume that the reference building measurement data is noise-free. As the measurement tolerance of equipment to measure temperatures and mass flows in buildings is usually given in a maximum percentage of the measured value, we add noise to each sample value  $s$  in the measurement data with

$$\tilde{s} = s + \alpha r s \text{ with } r \in [-1, \dots, 1], \quad (2.8)$$

where  $\tilde{s}$  denotes the new sample value,  $\alpha$  denotes the desired noise level (accuracy of the measurement equipment) and  $r$  is random, independent (for different samples) and

identically distributed uniformly between -1 and 1. For  $\alpha$ , values of 0.1, 0.2 and 0.3 are considered.

**Data preparation for varying data set size:** In the base case, the first 70% of the data set is used for training of the ANN and the remaining 30% for validation of the model. In order to ensure comparability between the different test cases, we keep the validation set the same throughout. To simulate limited availability of measurement data, the training set is shrunk from 41 weeks to one week by removing samples from the beginning, such that there is no temporal gap between the training and the validation set.

**Use of weather forecasts:** To investigate the effect of the quality of the weather forecast on the heat demand prediction accuracy, we compare prediction results obtained with measured weather data, which would represent a perfect forecast, to results obtained from a weather forecast. For the weather forecast, forecasts of the ambient temperature by MeteoSwiss, the Swiss national weather service, are used. Linear interpolation is used to generate 15 minute forecast samples from the forecast with hourly resolution.

## 2.6.2 Results and Discussion

### Influence of sampling time




Figure 2.17: Variance of the coefficient of determination for sampling times 15 minutes, 30 minutes, 60 minutes, 120 minutes; for (a) corrected and (b) uncorrected forecasts, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median. Note the different scale of the y-axes.

Figure 2.17 shows the prediction performance of the ANN with varying sampling rate of 15 minutes, 30 minutes, 60 minutes and 120 minutes for uncorrected and corrected (with both correction methods) demand forecasts. The forecasts for 30, 60 and 120 minutes were up-sampled by interpolation and validated against the test set for 15 minutes. In the uncorrected demand forecasts, a decreasing sampling time leads to lowering of the average coefficient of determination and also to a growing variance. A sampling time of 120 minutes does not lead to meaningful forecasts any more, which indicates that the Nyquist sampling rate is not respected any more. This effect could shift to higher or lower sampling times depending on the building inertia. In the corrected demand forecasts, the same trend occurs, but the negative effects of bigger sampling times are lowered. The variance is significantly decreased in all cases and the relative variance reduction increases with increasing sampling time (factors 4.9, 12.9, 14.5, 14.9 for sampling times 15, 30, 60, 120 subsequently). The average  $R^2$  is significantly raised by the correction methods for all sampling times. With an average  $R^2$  of 0.845 and an interquartile range of 0.008, a corrected forecast with a sampling time of 30 minutes leads to better results than an uncorrected forecast with 15-minute sampling time ( $R^2 = 0.823$  and  $IQR = 0.043$ ).

### Influence of noise level




Figure 2.18: Variance of the coefficient of determination for noise levels 0, 10%, 20%, 30%; for (a) corrected and (b) uncorrected forecasts, with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median. Note the different scale of the y-axes.

The results in Figure 2.18 show that the noise level has a noticeable influence on the




Figure 2.19: Forecast trajectory for test set with both correction methods applied with single week of offline training. The blue line denotes the true heating demand, the green line is the uncorrected forecast, and the orange line is the corrected forecast with both correction methods applied.

variance in the uncorrected case. However, the average  $R^2$  is lowered subsequently for higher noise levels. When online correction methods are applied, the variance appears to be constant for all noise levels. A noise level of 10% slightly increases the prediction performance for the corrected forecasts (average  $R^2$  of 0.892 vs. 0.885). Although this might appear surprising, adding noise to data is a common way to avoid overfitting in ANN (Zur et al., 2009). Further increase of the noise level reduces the prediction performance. By comparison with the results for varying sampling time, it can be seen that the performance is less sensitive to the noise level than to the sampling time.

### Influence of training data amount

Figure 2.19 shows the forecast trajectory of an uncorrected ANN and a corrected ANN with a single week of historical data for offline training. While the uncorrected forecast appears to have learned daily demand fluctuations, it does not catch the influence of the ambient conditions. The corrected forecast performs well from the start. It predicts daily demand fluctuations and shows dependency on the ambient conditions. After one month of operation, the performance does not differ from a model trained with 41 weeks of historical data. The achieved  $R^2$  of the shown corrected forecast is 0.860 (the average for 41 weeks of training is 0.885). The shown trajectory is not a “best case” example and the performance is repeatable.

### Influence of weather forecast

Figure 2.20 shows the prediction performance of the networks for (a) uncorrected demand forecasts and (b) corrected demand forecasts with different inputs used in the training and testing of the ANN: measured ambient temperature used both for training and

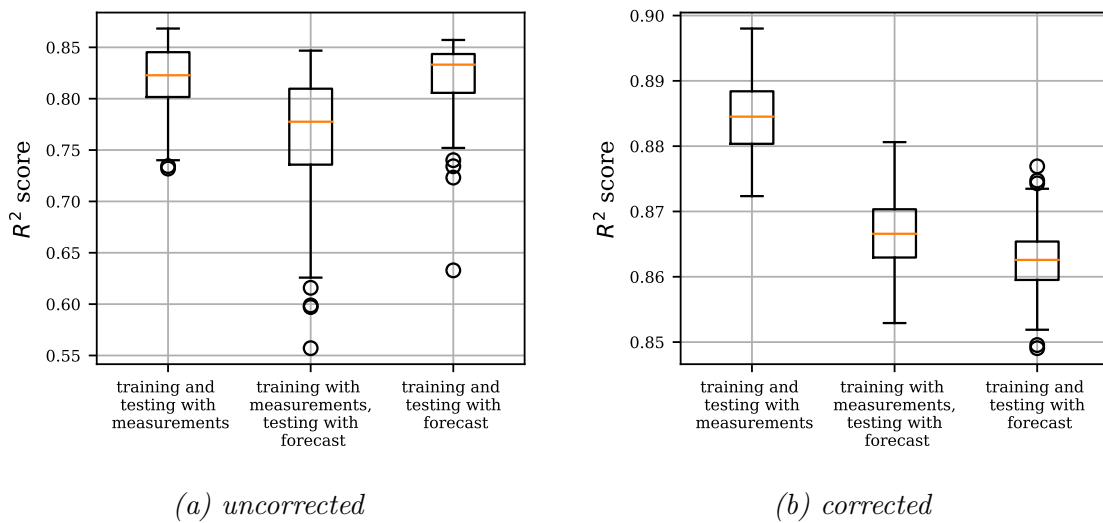


Figure 2.20: Variance of the coefficient of determination for different uses of weather forecasts with (a) uncorrected demand forecasts, and (b) corrected forecasts. Boxes describe the interquartile range, whiskers denote 1.5 times the interquartile range, circles indicate outliers and the orange line indicates the median. Note the different scale of the y-axes.

testing (perfect weather forecast), weather forecasts used for testing a model trained on measured ambient temperature, and weather forecasts used for both training and testing. The variance of the uncorrected ANN increases significantly when weather forecasts are used for forecasting but not for training. This is not the case if forecasts are used for both training and testing. However, the correction methods bring the variance to a constant level as can be seen in (b). The average  $R^2$  reduces from 0.885 to 0.866 and 0.863 when online correction methods are used. The comparison of networks that use ambient temperature measurements for training with those that use weather forecasts shows that the former appears to offer no significant benefit over the latter.

## 2.7 Conclusion

Artificial Neural Networks can be used for data-driven demand forecasting in the building and district energy domain. However, their prediction accuracy has high variance and depends on network parameters that are commonly randomly initialized. Here, we have introduced two simple forecast correction methods to significantly reduce this variance without using computationally intensive ensemble methods. The correction methods are based on the autocorrelation of the forecasting error and on online learning.

We have conducted three case studies to demonstrate the capabilities of the approach with respect to prediction performance, generalizing capability, and training data sensi-

tivity. With the help of the first case study of a complex building with district energy system characteristics, we have demonstrated that the methods significantly reduce variance in the prediction performance and also improve prediction accuracy and model bias for a sub-hourly day-ahead heat demand forecasting task. Furthermore, we have shown that ANN with both correction methods outperform other grey-box and black-box forecasting methods in the case study.

In a second case study we have validated the transferrability of the corrections methods and ANN architecture to other building types by applying the approach to three additional buildings. The results have shown that ANN with the named correction methods perform well when compared to fitted 5R3C resistor-capacitor building models. We have also demonstrated that the methods improve accuracy and reduce the variance in all cases, which improves confidence in using ANN in the frame of control tasks.

In a third case study, the performance sensitivity with respect to sampling time, noise level, amount of data, and accuracy of weather forecasts was investigated. It was found that the sampling time of the inputs has significant impact. An input noise level of 10% increases prediction performance when online correction methods are used, whereas higher noise levels lead to reduced performance. Furthermore, one week of historical data suffices to train reliable ANN for forecasting in the demonstrated case. Moreover, using ambient temperature measurements for training of the ANN appears to offer little benefit compared to using weather forecasts directly for training when online correction methods are applied.

The case studies give an indication that the developed methods are valid for more than one specific building and can potentially be generalized. However, they do not provide a formal proof. The confidence in the method is strengthened by the consistency of the results. The generated heating demand forecasts will be used as an input for a Robust Model Predictive Control scheme to provide frequency reserves to electrical grids based on demand-side management in buildings and districts in the next Chapter 3 and in Chapter 5.





# Robust MPC for frequency regulation with heat pumps

Accurate demand forecasting methods, as presented in Chapter 2, are a key ingredient to control schemes for demand-side management with buildings and districts. With the increased amount of volatile renewable energy sources connected to the electricity grid, and the phase-out of fossil fuel based power plants, there is an increased need for frequency regulation. On the demand side, frequency regulation services can be offered by buildings or districts that are equipped with electric heating or cooling systems by exploiting their thermal inertia. Existing approaches for tapping into this potential typically rely on dynamic building models, which in practice can be challenging to obtain and maintain. Moreover, actively controlling individual buildings in a district setting requires extensive control infrastructure, which might not be available in legacy configurations, or may cause privacy concerns for occupants. To address these issues, in this chapter, we instead exploit the thermal inertia of buffer storage for reserves, reducing the building models to demand forecasts. The developed control scheme paves the way for Chapter 5, where also building dynamics will be included. By combining a control scheme based on Robust Model Predictive Control with affine policies, and heating demand forecasting based on Artificial Neural Networks (ANN) with online correction methods, as introduced in Chapter 2, we offer frequency regulation reserves and maintain user comfort with a district system comprising a central heat pump and buffer storage. While the robust approach ensures occupant comfort, the use of affine policies reduces the effect of disturbance uncertainty on the system state. In a three-day experiment with a real district-like building energy system, we demonstrate that the scheme can offer reserves in a variety of conditions and track a regulation signal with high accuracy while meeting the heating demand of the connected buildings. Of the consumed electricity, 13.4% is flexible. In additional numerical studies, we demonstrate that using affine policies significantly decreases the cost function and increases the amount of offered reserves and we investigate the suboptimality in comparison to an omniscient control system.

## 3.1 Introduction

The amount of renewable energy sources in the electricity grid is continuously increasing. As many of these sources are highly volatile, there is a growing need for frequency regulation (Johnson et al., 2019). A common strategy for frequency regulation is the deployment of fast-reacting power plants, for example gas or hydro-power; an emerging strategy is the use of storage technologies, for example batteries. Besides such regulation on the supply side of the grid, frequency regulation on the demand side is possible through manipulation of controllable loads. This concept falls under the category of demand-side management (Gelazanskas and Gamage, 2014).

Possible candidates for demand-side management are buildings equipped with electric heating or cooling systems, such as heating, ventilation and air conditioning (HVAC) units, electric heaters and heat pumps (Hao et al., 2015), or entire district heating systems with electric heat sources. Due to the thermal inertia of buildings, they are to an extent flexible when it comes to their heating and cooling requirements, hence their electricity consumption. By shifting their consumption in time they can therefore influence the grid frequency (Fischer and Madani, 2017).

However, shifting electricity consumption can have an impact on occupant comfort as heating and cooling energy might not be available at the exact time when it is needed. There are different strategies to mitigate this impact. Rominger et al. (2018); Romero Rodríguez et al. (2019) and Maasoumy et al. (2013) develop and test control strategies for frequency regulation with heat pumps and HVAC units without explicitly enforcing comfort constraints and check only a-posteriori whether these were violated or not. Zhao et al. (2013) and Wang et al. (2020) use heuristics based on weather forecasts and occupancy to limit the offered frequency reserve capacity to enforce comfort constraints. Other authors use dynamic building models to exactly determine the influence of changed heating and cooling supply on room temperatures when providing reserves with heat pumps (Kim et al., 2015) or air handling units (Hao et al., 2013; Lin et al., 2013; Olama et al., 2018).

Combined with optimization in the frame of Model Predictive Control (MPC), such dynamic models can be used to maximize the offered frequency reserves while maintaining comfort constraints (Rastegarpour et al., 2020a,b). Zhang et al. (2017) and Vrettos et al. (2016) develop Robust MPC schemes to provide day-ahead reserves for frequency regulation with commercial buildings and HVAC systems. Robust schemes ensure occupant comfort in the face of uncertainty in the regulation signal from the transmission system operator (TSO). In (Vrettos et al., 2018b,a) this approach is further developed and tested in a case study on a real small air-conditioned building. In (Oldewurtel et al., 2013), MPC is used to provide ancillary services with a real building, and simulation studies are conducted to investigate a setup of an aggregation of office buildings.

As buildings and related components (such as buffer storage) have integrator dynamics, uncertain disturbances lead to uncertainty accumulation in the system state and thus limit the electrical reserves offered. This is the case because the Robust MPC uses much of the available flexibility to account for the uncertainty instead of exploiting it for reserves. Feedback policies are a known tool in Stochastic and Robust MPC to mitigate this problem (Goulart et al., 2006), as they allow the controller to plan to react on uncertainty before it is revealed. However, to the best of the authors' knowledge, they have not yet been applied in practice to reserve provision with buildings and districts.

Obtaining dynamic models of the thermal dynamics of buildings can be challenging, as will be discussed in Chapter 4. Furthermore, most of today's buildings are not yet equipped with the necessary hardware for predictive room temperature control. Moreover, if buildings are to be controlled from an external entity in the context of reserve provision, for example by an aggregator (Oldewurtel et al., 2013), occupants might have privacy concerns related to room temperature and other live measurements, and could therefore prefer to not actively participate. To also account for these cases, alternative strategies for reserve provision schemes that do not necessarily rely on exploiting the thermal inertia of buildings are necessary.

Here, we take a step in this direction by decoupling heating/cooling demand and heating/cooling supply of the system. This is achieved by placing a buffer storage, e.g. a water tank, between supply and demand. In this case, the thermal inertia of the buffer storage allows flexibility in heating and cooling energy production without the necessity to exploit the inertia of the connected building itself. The comfort of the occupants is ensured by a separate (potentially unknown) lower level controller in the building as long as the buffer storage has a sufficiently high temperature. The heating/cooling system and storage can in this case be modeled with first principles, which is tractable from an economic point of view as these are mass-produced products. The demand of the building can be modeled with a forecasting method, such as the method introduced in Chapter 2.

The above named configuration of a central heating/cooling device connected to buffer storage is not only common for many commercial buildings, but also for entire district heating systems (Hennessy et al., 2019). Here, central heat pumps and buffer storages are controllable by the system operator, while connected buildings might not be controllable (e.g. due to privacy concerns or hardware/model limitations). These can be forecast, but not directly influenced. Using the thermal buffer presents an option to offer electricity reserves even if the building's thermal mass cannot be controlled, due to one or more of the reasons given above.

There is a growing interest in the potential (Lund et al., 2018; Ivanova et al., 2019) of providing ancillary services in district heating systems, and there are attempts to quantify flexibility potentials (Xu et al., 2020) and to model individual components for this purpose (Pagh Nielsen and Sørensen, 2021). However, review studies (Vandermeulen

et al., 2018; Hennessy et al., 2019) indicate that there is little work on the operation (Terreros et al., 2020; Salpakari et al., 2016; Li et al., 2016) of such systems with reserve schemes and no implementation and validation on a real system in the literature. We feel that enabling deployment at the district level is an important step because, with the current structure of ancillary markets, individual buildings can only enter through massive aggregation (Geidl et al., 2017; Koch et al., 2011; Borsche et al., 2014; Oldewurtel et al., 2013); districts, however, may be able to enter individually, or in small clusters, if minimal power requirements are reached. Moreover, deploying the methods in districts instead of individual buildings can improve performance because of increased demand predictability thanks to averaging (Bünning et al., 2020a) and because of synergistic effects, for example between residential and commercial buildings (Darivianakis et al., 2017).

However, for a practical application of such a reserve scheme, the combination of advanced methods is essential. High-accuracy demand forecasting methods are needed to reduce the uncertainty in the disturbance of the buffer storage, and as a result the uncertainty in the system state. Robust MPC (Vrettos et al., 2016) is needed to ensure occupant comfort in the face of uncertainty of the demand forecasts, of the frequency regulation signal and of other disturbances. Moreover, to maximize the reserves offered, feedback policies (Goulart et al., 2006) are required, to minimize the effect of disturbance uncertainty on the system state, and as a result to better exploit the available storage compared to open-loop MPC. This is especially important in the considered configuration, as the thermal inertia of buffer storage is considerably smaller compared to the inertia of the thermal mass of buildings.

### **3.1.1 Contribution**

In this Chapter, we therefore present a combination of a three-level Robust MPC (Vrettos et al., 2018a) for frequency regulation approach with the forecasting methods presented in Chapter 2, to offer day-ahead frequency regulation reserves with a system comprising a central ground-source heat pump and water buffer storage that meet the heating demand of a building or group of buildings. The robust formulation is combined with affine policies, as discussed in (Warrington et al., 2012) for reserve provision in power systems, which allows us to increase the reserves offered compared to standard open-loop MPC. While the individual methods have been presented in previous work, their combination is required to achieve a practical implementation to enable the deployment of ancillary services to the level of districts with central heat supply. We also document what we believe to be the first implementation demonstrating this approach in experiments: We validate the methods on a physical system, the NEST demonstrator in Switzerland, a “vertical neighbourhood” connected by a district heating system (Richner et al., 2017). We show that the approach is able to offer a substantial amount of regulation reserves in

a variety of ambient conditions and on a configuration, where the storage capacity is only a small fraction of the daily heating demand. We also demonstrate that the approach ensures good regulation signal tracking performance with a variable speed heat pump. Furthermore, we investigate optimality properties of the MPC solutions and investigate the effect of the use of affine policies in two numerical experiments.

The presented scheme paves the way for Chapter 5, where a subset of the connected buildings will be considered as uncontrollable (i.e., are treated as presented in this chapter), and other connected buildings are considered as controllable, i.e., their thermal mass can be exploited. This is after data-driven models for building thermal dynamics are introduced in Chapter 4.

### 3.1.2 Structure

The remainder of the Chapter is structured as follows. In Section 3.2, we introduce the reserve provision scheme and the system under consideration. In Section 3.3, we discuss the models for heat pump and storage and re-visit the prediction models with correction methods for the heating demand of buildings from Chapter 2. We also describe the Robust MPC based control scheme and the use of affine policies. In Section 3.4, we present the experimental case study and its results. In Section 3.5, we describe the numerical case studies and discuss the suboptimality of the presented approach. We discuss limitations and future research directions in Section 3.6 and conclude the chapter in Section 3.7.

## 3.2 Problem statement

### 3.2.1 Reserve provision scheme

We assume a day-ahead planning frequency regulation reserve scheme, as used by the regulation product RegD, offered by the U.S. transmission system operator PJM (*Pennsylvania, New Jersey, and Maryland*). We use the PJM reserve market to have a concrete example. However, the method is not limited to this scheme and can be easily adapted to other ancillary market conditions. In the considered scheme, the reserve provider communicates an offer  $r \in \mathbb{R}_+^{96}$  of symmetric reserves to the TSO at midnight. The offer is made in 15-minute intervals and is fixed for the next 24 hours. During the next day, when the offered reserves are due, the reserve provider can change its base consumption  $u_k^0$  every timestep  $k$  (i.e. every 15 minutes). It should then track the electrical load

$$u_k(\tau) = u_k^0 + w(\tau)r_k, \quad (3.1)$$

where  $w(\tau) \in [-1, 1]$  denotes the regulation signal which is updated every 2 seconds by

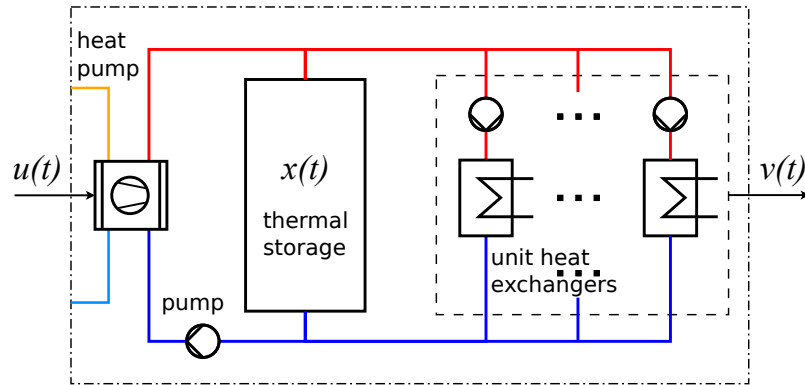


Figure 3.1: Schematic of the system under consideration with heat pump, water storage tank and heat exchangers for individual apartments.

the TSO, and  $r_k$  denotes the  $k^{\text{th}}$  element of the list of offered reserves  $r$ . The base load  $u_k^0$  can be updated every 15 minutes,  $r_k$  is time varying in 15 minute intervals but fixed for the day, while  $w(\tau)$  changes every 2 seconds. Thus,  $u_k(\tau)$  also changes and is sampled every 2 seconds.

The tracking performance is judged by a *composite performance score* monitored by the TSO, which consists of an *accuracy score* which measures the correlation between the reserve signal and the system response, a *delay score* which measures the time delay between reserve signal and system response, and a *precision score* which measures the error between reserve signal and system response (PJM, 2019).

Our scheme does not consider time-varying energy prices and peak pricing for  $u_k^0$ , as these aspects have been treated in the literature already, and their handling would dilute the scope of this work. We therefore refer the interested reader to (Zhang et al., 2015; Sundström et al., 2017; Ma et al., 2014).

### 3.2.2 System under consideration

We consider the heating system for reserve provision shown in Figure 3.1. It consists of a vapour compression cycle heat pump, which is depicted on the left, and a water storage tank, which is depicted in the middle. The heat pump draws cold water from the bottom of the storage tank with the help of a pump, warms up the water by transferring heat from the refrigerant to the water inside the condenser, and feeds it back into the top of the storage. By varying the heat pump's electrical consumption, frequency regulation can be offered. On the right, individual pumps draw warm water from the top of the storage tank and pass it through heat exchangers, which supply individual apartments of a building or individual buildings of a district with heat. The cold water is returned to the bottom of the tank.

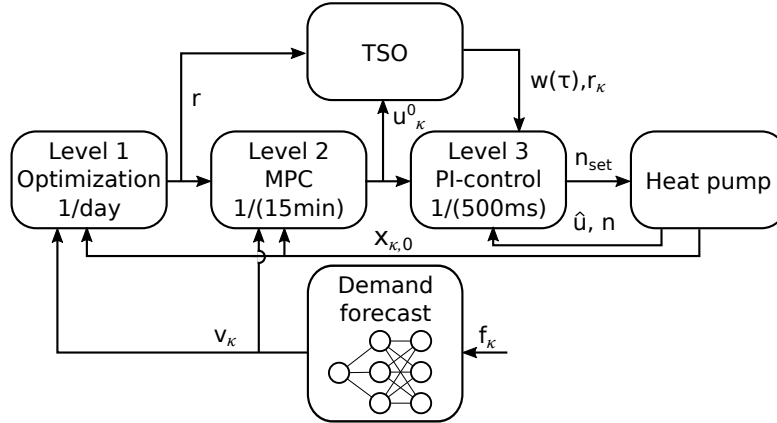


Figure 3.2: Hierarchical control scheme with three controller levels.

### 3.2.3 Control scheme

To provide day-ahead frequency regulation services, we use the three-level scheme shown in Figure 3.2. The scheme is an extension of (Vrettos et al., 2018a) and was preliminarily tested in a software-in-the-loop experiment in (Bünning et al., 2020c). Level 1 is an optimization problem solved once per day, to determine the offered reserves  $r$ , which are communicated to the TSO and fixed for the day ahead. Level 2 is an MPC controller, which re-optimizes the base load  $u_{\kappa}^0$  of the heat pump every fifteen minutes and reacts to unforeseen disturbances.<sup>1</sup> Level 3 is a lower-level feedback controller which tracks the regulated load with a fast sampling time of 0.5 seconds. Levels 1 and 2 use a heating demand forecast as an input, for example the method developed in Chapter 2. The required models for levels 1 and 2, along with the detailed control scheme are presented in the following section.

## 3.3 Methodology

### 3.3.1 Models

The heat pump and water storage are modeled with first principles (physics based), while the heating demand of the building is modeled with the help of online corrected ANN. This is done as heat pumps and storage tanks are mass-produced industrial products for which first principles models are relatively easy to develop, while buildings are generally different from each another and thus modeling the building demand with first principles would require significant effort for each building.

<sup>1</sup>The index  $\kappa \in [1, 96]$  denotes the discrete time index, e.g.  $\kappa = 1$  for midnight and  $\kappa = 2$  for 00.15 a.m., while  $k \in [1, N]$  denotes the index in the optimizations: for example,  $u_{3,8}^0$  is the eighth element of the heat pump base consumption in the optimization conducted at time  $\kappa = 3$  (00.30 a.m.).

### Heat pump and storage model

The heat pump, depicted on the left of Figure 3.1, generates high temperature heat  $u_{\text{th}}(t)$  by using electricity  $u(t)$  and ambient heat at a lower temperature level. Here,  $t$  denotes continuous time. The conversion efficiency between electrical energy and high temperature thermal energy is assumed to be described with the coefficient of performance (COP)  $\alpha_{\text{COP}}$ :

$$u_{\text{th}}(t) = \alpha_{\text{COP}} u(t) + e(t). \quad (3.2)$$

To keep the optimization problem linear, a constant  $\alpha_{\text{COP}}$  is a reasonable assumption for ground-sourced heat pumps: as the temperature of the ground changes slowly and the heating supply temperature on the building side is usually fixed for a day (following a *heating curve* dependent on the ambient temperature of the previous day), the COP can be updated on the basis of previous days' measurements for example, and kept constant for a day. For air-sourced heat pumps,  $\alpha_{\text{COP}}$  can be modeled as a function of the ambient temperature, which in practice can be obtained from the weather forecast.

The error  $e(t)$  captures potential additional thermal disturbances. Note that equation (3.2) holds for any  $u_{\text{th}}$  and  $u$ , thus also for the discrete time instants used in the reserve scheme of (3.1).

Neglecting thermal losses, because they are slow and small compared to the charging and discharging of the tank, the average temperature  $x(t)$  of the storage tank in Figure 3.1 is described by the energy balance

$$m c_p \frac{dx(t)}{dt} = u_{\text{th}}(t) - v(t) + \delta(t), \quad (3.3)$$

where  $m$  and  $c_p$  denote mass and isobaric specific heat capacity of the water respectively,  $v(t)$  denotes the heating demand of the building, and  $\delta(t)$  denotes the error between the forecast and the actual heating demand. Like the error  $e(t)$  in (3.2),  $\delta(t)$  will be modeled as a box-constrained uncertainty set for the robust optimization. Allowing mixing of different water layers in the storage, but assuming no swapping of temperature layers<sup>2</sup>, the average temperature constitutes a lower bound for the water temperature in the top layer and an upper bound for the temperature in the lowest layer, which is sufficient for our control purpose. Moreover, model inaccuracies compared to a stratified tank model can also be captured by  $\delta(t)$ . Inserting equations (3.1) and (3.2) into equation (3.3) gives rise to the full linear description of the storage temperature:

$$m c_p \frac{dx(t)}{dt} = \alpha_{\text{COP}} (u_k^0 + w(\tau)r_k) + e(t) - v(t) + \delta(t). \quad (3.4)$$

---

<sup>2</sup>The results of the experiment will confirm that this is a reasonable assumption.



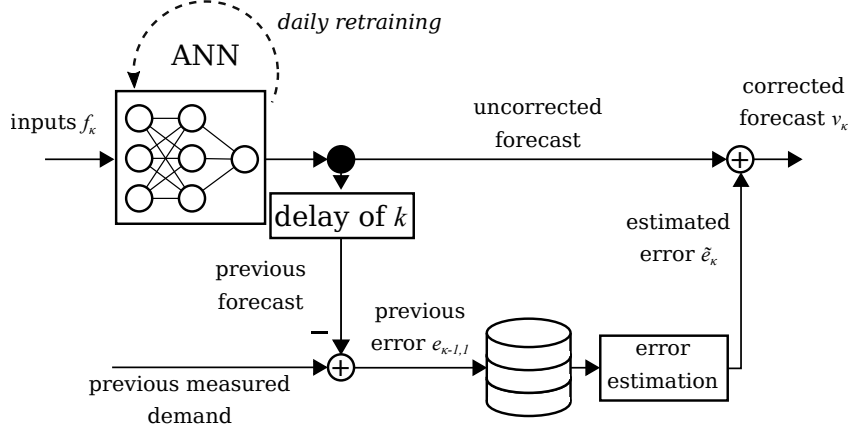


Figure 3.3: Forecast correction based on error-autocorrelation and online learning.

### Building energy demand model

The ANN forecasting approach with online correction methods for forecasting heating demands of buildings and districts has been presented in Chapter 2. For the sake of completeness, we briefly revisit the methods here and adapt them to the addressed forecasting task. For the purposes of frequency reserve provision, a heating demand forecast for a building for the next 24 hours is made starting at midnight and afterwards every 15 minutes until the end of the day. The forecasting horizon thus decreases by 15 minutes with every forecast. Both training and validation data are assumed to be sampled at 15 minute time steps. The forecast is made with a feed-forward ANN, with inputs related to ambient conditions and time features.

Two correction methods are applied in the online phase of the forecasting task (Figure 3.3). The first correction method is based on the forecasting **error-autocorrelation**. The error  $\tilde{e}$  of the forecast conducted at the current time  $\kappa$  for forecasting interval  $k$  is estimated with

$$\tilde{e}_{\kappa,k} = e_{\kappa-1,1} R_{ee}(k, \mathcal{E}), \quad (3.5)$$

where,  $e_{\kappa-1,1}$  denotes the difference between the first (15-minute) element of the last conducted forecast (at time  $\kappa - 1$ ) and the actual measured heating demand.  $R_{ee}(l, \mathcal{E})$  is the autocorrelation of the forecasting error, which is dependent on a time-lag  $l$  and the set of all past forecasting errors  $\mathcal{E}$ , as described in Equation (2.2).

The correction procedure is illustrated in Figure 3.3. A forecast is made at time  $\kappa$  based on the inputs  $f_\kappa$ . The previous forecast from time  $\kappa - 1$  is compared to the actual measured heating demand, giving rise to  $e_{\kappa-1,1}$ . With all previously measured errors  $\mathcal{E}$ , stored in a database,  $\tilde{e}_\kappa$  can be calculated on the basis of (3.5). Adding  $\tilde{e}_\kappa$  to the uncorrected forecast gives rise to the corrected forecast  $v_\kappa$ , which will later be used as

an input to the control scheme.

The second forecasting correction method is based on **online learning**: Instead of only training the ANN on a training set offline and using the ANN for predictions online, the ANN is retrained online every 24 hours on the basis of the data gathered during the previous day. By doing this, changes to the building that persist for longer than one day can be captured; such changes could include changing the set point of a thermostat for example. While conventional wisdom suggests that ANN require large data sets for training, due to the correction methods, the presented approach can already reach high prediction performance with a single week of training data as shown in Chapter 2, Section 2.6.

### 3.3.2 Controller design

The models and demand forecast developed in Section 3.3.1 are used in the three-level control scheme depicted in Figure 3.2. Level 1 solves a robust optimization problem once every 24 hours at midnight. Based on the current storage tank temperature  $x_{\kappa,0}$  and the heating demand forecast of the building  $v_{\kappa}$  it determines the reserves  $r$  to be offered and fixed in 15 minute intervals,  $r_k$ , over the next 24 hours. Level 2 solves an optimization problem similar to the one in Level 1 every 15 minutes during the day, with a shrinking horizon, from the current time to midnight. In this optimization problem, the values of the reserves  $r$  for the rest of the day are already known, because they have been fixed by Level 1. Further inputs to this problem are the updated tank temperature measurement  $x_{\kappa,0}$  and the updated heating demand prediction  $v_{\kappa}$ . The outputs of Level 2 are the nominal heat pump electrical power set points  $u_{\kappa}^0$  for each 15 minute interval for the remainder of the day, of which the first one,  $u_{\kappa,1}^0$ , is passed on to Level 3. Level 3 is a Proportional-Integral controller that controls the relative rotational speed  $n$  of the heat pump's compressor to track the regulated heat pump's electricity consumption  $u_{\kappa}(\tau)$ .

#### Level 1

For Level 1, equation (3.3) can be rearranged to the continuous state space system

$$\begin{aligned} \frac{dx(t)}{dt} &= \frac{1}{m c_p} (u_{\text{th}}(t) - v(t) + \delta(t)) \\ &= A_c x + B_c (u_{\text{th}}(t) - v(t) + \delta(t)), \end{aligned} \quad (3.6)$$

with  $A_c$  and  $B_c$  denoting the continuous state and input matrix respectively. Discretizing in time gives the discrete state space system

$$x_{k+1} = \tilde{A}x_k + \tilde{B}(u_k - v_k + \delta_k), \quad (3.7)$$

where  $\tilde{A}$  and  $\tilde{B}$  are the discrete time state and input matrices respectively. By redefining  $x := [x_1, \dots, x_N]^\top \in \mathbb{R}^N$ ,  $u_{\text{th}} := [u_{\text{th},1}, \dots, u_{\text{th},N}]^\top \in \mathbb{R}^N$ ,  $v := [v_1, \dots, v_N]^\top \in \mathbb{R}^N$ ,  $\delta := [\delta_1, \dots, \delta_N]^\top \in \mathbb{R}^N$ , with  $\top$  denoting a transpose, and  $N$  denoting the prediction horizon, we can describe the discrete state trajectory by

$$x = Ax_0 + B(u_{\text{th}} - v + \delta), \quad (3.8)$$

where  $x_0$  is the initial state of the system, and  $A$  and  $B$  are defined appropriately (Borrelli et al., 2017). By vectorizing all remaining variables and adding equation (3.2) to the formulation, the robust optimization problem in terms of the offered reserves  $r$ , the nominal heat pump electrical set points  $u^0$ , and the heat pump on/off condition  $z$ , can be written as

$$\min_{r,x,u^0,u_{\text{th}},\epsilon,z} \quad f^{\text{el}\top} u^0 - f^{r\top} r + \lambda^\top \epsilon \quad (3.9a)$$

$$\text{subject to} \quad x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (3.9b)$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + w \odot r), \quad (3.9c)$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (3.9d)$$

$$zU_{\min} \leq u^0 + w \odot r \leq zU_{\max}, \quad (3.9e)$$

$$z \in \mathbb{Z}_2^N, \quad (3.9f)$$

$$\epsilon \geq 0, \quad (3.9g)$$

$$\forall w \in W, \forall \delta \in \Delta, \forall e \in E. \quad (3.9h)$$

Here,  $f^{\text{el}}$  and  $f^r$  denote costs for electricity and benefits for offered reserves respectively.  $X_{\min}$  and  $X_{\max}$  describe temperature limits for the storage tank, defined by the lowest possible operating temperature for floor heating, which ensures the indoor thermal comfort, and the highest supply temperature of the heat pump. The slack variable  $\epsilon \in \mathbb{R}^N$  ensures feasibility with respect to the storage temperature constraint and  $\lambda$  denotes the associated cost. The lower and upper electrical capacity limits of the heat pump are described by  $U_{\min}$  and  $U_{\max}$ , and  $z \in \mathbb{Z}_2^N$  is a binary variable that determines if the heat pump is switched on or off. The symbol  $\odot$  denotes the operator for element-wise multiplication. All constraints have to hold for all realizations of uncertainties  $w \in W, \delta \in \Delta, e \in E$ . Suitable sets for  $W$  can be found by analysing the regulation signal  $w(\tau)$ . Note that  $r$  will always be positive due to the formulation of the cost function. However, as  $W$  can contain negative regulation signals, the effective reserve can become negative.

Constraints (3.9c) and (3.9e) can be reformulated as linear constraints by making  $w$  a square diagonal matrix. While  $W$ ,  $\Delta$  and  $E$  generally allow any convex sets, for

box-constrained sets the robust optimization problem (3.9), which has to hold for the qualifier (3.9h), can be reformulated as a Mixed Integer Linear Program via *explicit maximization* (Löfberg, 2012). This reformulation is performed automatically by many modern optimization tools (Löfberg, 2004; Goh and Sim, 2011).

The binary variable  $z$  forces  $u^0$  and  $r$ , and thus  $u_{\text{th}}$ , to be zero if the electrical input to the heat pump does not exceed  $U_{\text{min}}$ . For the heating system this means that in case of low heating demand from the building, a hysteresis behaviour can be expected, where the heat pump changes between over-serving the demand and switching off. Potentially, there could be combinations of  $X_{\text{min}}$ ,  $X_{\text{max}}$ ,  $U_{\text{min}}$  and  $U_{\text{max}}$  where the heating demand could not be served, but this issue is captured by the slack variable  $\epsilon$ .

In the case of low storage capacities, i.e. large  $B$ , corresponding to low mass of water, low  $X_{\text{max}}$  or high  $X_{\text{min}}$ , or large uncertainty in  $W$ ,  $\Delta$ , and  $E$ , the offered reserves  $r$  may become very small or, without the slack variable  $\epsilon$ , the problem may even become infeasible. This is because the uncertainty induced in  $x$  by the action of  $w$ ,  $\delta$ , and  $e$  compounds along the horizon, as uncertainty at subsequent steps gets added to that of earlier steps through the “integrator” implicit in (3.9b) (see (3.3)). As a consequence, near the end of the horizon the uncertainty in  $x$  becomes large, leading to a violation of (3.9d). This growth in uncertainty traces its origins to the fact that (3.9) addresses Level 1 in the control hierarchy of Figure 3, but does not contain any information about the actions of the lower levels. In reality, Level 2 and Level 3 will be executed repeatedly within the horizon of (3.9), adjusting the decisions of Level 1 to account for information that has become available in the meantime. This introduces feedback to the process, that will in practice limit the growth of the uncertainty.

In stochastic programming, information about this “recourse” process can be introduced by optimising over causal feedback policies instead of a sequence of “open-loop” decisions fixed at the beginning of the horizon. In this case, the optimization problem for Level 1 encodes the fact that the system will react to uncertainties that are still unknown at the time Problem (3.9) is solved, but will be revealed at the time the decision is implemented. Unfortunately, as discussed by Warrington et al. (2012), optimizing over the set of all possible policies is intractable in general. To obtain a tractable optimization problem, one can restrict the classes of causal policies considered. A common choice in this respect is the class of affine disturbance policies (Ben-Tal et al., 2004; Goulart et al., 2006). For the uncertainties introduced by the regulation signal  $w$ , equation (3.9c) can be extended to

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + w \odot r + D_w w), \quad (3.10)$$

where  $D_w \in \mathbb{R}^{N \times N}$  is a strictly lower triangular matrix:

$$D_w := \begin{pmatrix} 0 & 0 & \cdots & 0 \\ [D_w]_{2,0} & 0 & \ddots & 0 \\ \vdots & \ddots & \ddots & 0 \\ [D_w]_{N,0} & \cdots & [D_w]_{N,N-1} & 0 \end{pmatrix}. \quad (3.11)$$

By making  $D_w$  a decision variable in the optimization problem, the uncertainty in  $u_{\text{th}}$  can be lowered, and thus also the uncertainty in  $x$ . Affine policies on the other uncertain variables  $\delta$  and  $e$  can also be defined. Because  $\delta$  and  $e$  appear together in (3.9b), a single lower triangular matrix can be used:

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + w \odot r + D_w w + D_{\delta,e}(\delta + e)). \quad (3.12)$$

We note that the regulation signal that takes values in the interval  $[-1, 1]$  is updated every 2 seconds, but the rest of the decision variables in (3.9) refer to quantities that are updated every 15 minutes. Therefore, when trying to meet the robust constraint (3.9d), the average value of  $w(\tau)$  over a 15 minute interval (denoted by  $\bar{w}$  below) is more relevant than the instantaneous value. By collecting historical data, a second uncertainty set on the average of  $w(\tau)$  can be created by integrating over 15-minute horizons and evaluating the distribution of these integrals (see Vrettos et al. (2018a)). As a result, the uncertainty set is decreased to  $\bar{W} \subset W$  for constraint (3.9c). As the instantaneous electrical consumption needs to remain within operational limits at all times,  $w \in W$  remains for constraint (3.9e).

The resulting optimization problem is

$$\min_{\substack{r, x, u^0, u_{\text{th}}, z, \\ \tilde{z}, D_w, D_{\delta, e}, \epsilon}} f^{\text{el}\top} u^0 - f^r \top r + \lambda \top \epsilon \quad (3.13\text{a})$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (3.13\text{b})$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + \bar{w} \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e)), \quad (3.13\text{c})$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (3.13\text{d})$$

$$zU_{\min} \leq u^0 + w \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e) \leq zU_{\max}, \quad (3.13\text{e})$$

$$\tilde{z}R_{\min} \leq r \leq \tilde{z}R_{\max}, \quad (3.13\text{f})$$

$$z, \tilde{z} \in \mathbb{Z}_2^N, \quad (3.13\text{g})$$

$$\epsilon \geq 0, \quad (3.13\text{h})$$

$$[D_w]_{i, j} = 0 \forall j \geq i, \quad (3.13\text{i})$$

$$[D_{\delta, e}]_{i, j} = 0 \forall j \geq i, \quad (3.13\text{j})$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta \in \Delta, \forall e \in E. \quad (3.13\text{k})$$

The heat pump capacity constraint, now (3.13e), is adapted to ensure feasibility under the chosen policies. Note that,  $[D_{\delta, e}]_{k, j}$  and  $[D_w]_{k, j}$  (as well as  $u_k^0$  and  $r_k$ ) will be zero whenever  $z_k = 0$ . Moreover, as preliminary experiments suggested that small reserves  $r$  lead to weak tracking performance (and low performance scores) because of large relative errors, a second binary variable  $\tilde{z}$  was added to impose a lower limit on  $r$  through constraint (3.13f).

## Level 2

Controller Level 2 is an MPC scheme with shrinking horizon. The horizon is shrinking, because  $r$  is only fixed by Level 1 until the end of the day. This choice could potentially lead to storage depletion towards the end of the day if no reserves are offered during these intervals. However, as we use a slack variable on the state constraints, there are no issues with recursive feasibility.<sup>3</sup> At most, storage depletion will lead to low reserves being offered in the first intervals of the following day. Alternative formulations include a receding horizon scheme, which allows Level 2 to choose those reserves  $r$  at the end of the horizon that are not yet fixed by Level 1. However, this significantly increases the computational complexity. Alternatively, the reserves that are not fixed can be padded with zeros, as will be discussed in Chapter 5.

Level 2 can update  $u^0$ , depending on new measurements of initial conditions  $x_0$  and demand forecasts  $v$ , which are updated every 15 minutes. It uses an optimization problem

---

<sup>3</sup>Moreover, as the system is stable and in practice often operated with hysteresis control, even without slack variables, recursive feasibility is ensured.

similar to that of Level 1. The only difference is that  $r$  is now fixed, leading to

$$\min_{\substack{x, u^0, u_{\text{th}}, z, \\ D_w, D_{\delta, e}, \epsilon}} f^{\text{el}\top} u^0 + \lambda^\top \epsilon \quad (3.14\text{a})$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (3.14\text{b})$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + \bar{w} \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e)), \quad (3.14\text{c})$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (3.14\text{d})$$

$$zU_{\min} \leq u^0 + w \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e) \leq zU_{\max}, \quad (3.14\text{e})$$

$$z \in \mathbb{Z}_2^N, \quad (3.14\text{f})$$

$$\epsilon \geq 0, \quad (3.14\text{g})$$

$$[D_w]_{i, j} = 0 \quad \forall j \geq i, \quad (3.14\text{h})$$

$$[D_{\delta, e}]_{i, j} = 0 \quad \forall j \geq i, \quad (3.14\text{i})$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta \in \Delta, \forall e \in E. \quad (3.14\text{j})$$

Other than  $r$ , no information from Level 1 is carried over to Level 2. The purpose of Level 2 is to introduce feedback, i.e. change the planned control inputs due to newly available information. The optimization schemes in Levels 1 and 2 do not consider electricity peak pricing and time-varying energy prices. However, they can be easily included in the schemes as described in (Zhang et al., 2015; Sundström et al., 2017; Ma et al., 2014).

### Level 3

Level 3 is a discrete Proportional-Integral feedback controller, with proportional gain  $k_p$  and integral gain  $k_i$ , to track equation (3.1) with the heat pump. The controller output is the set point for the relative rotational compressor speed of the heat pump  $n_{\text{set}}$ . The controller input is the heat pump's measured electrical load  $\hat{u}$ . An anti-windup scheme is used in case the heat pump reaches its compressor speed limitations. The integration block of the controller is also bypassed if the difference between the set compressor speed  $n_{\text{set}}$  and the measured compressor speed  $n$  exceeds a limit. This is done because heat pumps usually have up and down ramping limits.

## 3.4 Experimental case study

### 3.4.1 Configuration

We test the reserve scheme in a three-day experiment on a real system in the NEST building (see Figure 2.2 in Chapter 2) at Empa, Switzerland. As discussed in Chapter



Figure 3.4: Experimental facility with heat pump (left) and water storage tanks (right) (Bünning et al., 2020c).

Table 3.1: Parameters for controller Levels 1 and 2.

$N = 96,$	$\alpha_{\text{COP}} = 3.53,$	$\bar{A} = 1,$
$\lambda = 5,$	$W = [-1, 1],$	$\bar{B} = 0.0978 \frac{K}{kW},$
$f^{\text{el}} = 1,$	$\bar{W} = [-0.25, 0.25],$	$X_{\text{min}} = 28^\circ\text{C},$
$f^{\text{r}} = 1.5$	$U_{\text{min}} = 8.2\text{kW},$	$X_{\text{max}} = 38^\circ\text{C},$
$R_{\text{min}} = 0.4\text{kW}$	$U_{\text{max}} = 12.8\text{kW},$	$E \oplus \Delta = [-4.0, 4.0] \text{kW}$

2, the building consists of individual residential, office and multi-use units that can be added and removed from the building backbone, as well as permanent office and meeting rooms. The individual units are connected to a central heating system with a supply temperature of 38 °C and a return temperature of 28 °C via heat exchangers and are equipped with their independent control systems considered to be unknown in the experiment. Due to the independence of the connected units and their controllers, and the structure of the heating system, the building is often referred to as a “vertical district” (Empa, 2018) or “vertical neighborhood” (Richner et al., 2017) and is intended to serve as an emulator for studies on concepts considered for district heating systems.

The heating system (Figure 3.4) comprises a ground source heat pump, specifically the two-compressor model WP-WW-2NES 20.F4-2-1-S-P100 produced by Viessmann with a maximum thermal capacity of 100 kW, and a water buffer storage consisting of two 1100 litre Matica water tanks connected in series. Only the first compressor stage is used in the experiment, which leads to a rated maximum thermal capacity of 50 kW. The system resembles the configuration described in Section 3.2.

The control scheme configuration is as follows. The parameters for controller Levels 1 and 2 are shown in Table 3.1. The cost-function related parameters  $\lambda$ ,  $f^{\text{el}}$ , and  $f^{\text{r}}$  were chosen based on preliminary numerical studies to balance the trade-off between cost optimality and constraint violation. Compared to the preliminary experiments in



(Bünning et al., 2020c), we have increased the reserve benefit  $f^r$  to get richer  $r$  vectors to test the robustness of the controller. They are not directly related to a specific energy product and therefore considered dimensionless. The values for  $\alpha_{\text{COP}}$ ,  $U_{\text{min}}$ ,  $U_{\text{max}}$  and  $R_{\text{min}}$  were set on the basis of preliminary heat pump experiments. The limits for  $W$  are properties of the used regulation signal RegD by PJM. The uncertainty set  $\bar{W}$  can be determined by analyzing historical regulation signals (Vrettos et al., 2018a),  $\tilde{A}$  follows from the assumption of no thermal losses, and  $\tilde{B}$  is calculated on the basis of the tank volume and the specific heat capacity of water. In contrast to (Bünning et al., 2020c), where we chose the set boundaries based on historical measurement data from the building and the heat pump, we do not specify the uncertainty sets  $E$  (error from disturbances of the heat pump) and  $\Delta$  (error from demand forecast) separately, but instead define the Minkowski sum  $E \oplus \Delta$ , and shrink it compared to the original source. The initial value of the average tank temperature,  $x_0$ , is determined by taking a weighted average of six temperature sensor measurements at different heights within the storage tank. To reduce wear on the heat pump, we introduce an additional constraint to Level 1, and require  $z_k$  to be constant during each thirty minute interval.

The proportional and integral gains,  $k_p = 2.0$  and  $k_i = 0.4$ , of the PI controller in Level 3 were determined by first modeling and auto-tuning a first-order representation of the heat pump in Simulink<sup>®</sup>, and then manually adjusting the values after implementing the controller on the actual plant. The controller output limits for  $n_{\text{set}}$  are set to 20% and 50% of the relative compressor speed<sup>4</sup>. The controller sampling time is 500 milliseconds and the maximum allowed difference between controller output  $n_{\text{set}}$  and measured compressor speed  $n$  before anti-windup activates is set to 2%. Moreover, as the heat pump is only fully controllable five minutes after switching it on, it is switched on five minutes early in the case that reserves are offered for the next 15-minute interval.

For the implementation of Levels 1 and 2 we use Matlab<sup>®</sup>. As all uncertain variables are box-constrained, the optimization problems become MILPs. These are written with YALMIP (Löfberg, 2012), which automatically derives robust counterparts, and are solved with CPLEX<sup>®</sup> 12.9. Each optimization is started five minutes before the decision is implemented, limiting the solver time to five minutes. This time is enough to solve the problem close to optimality in all cases.<sup>5</sup> The best feasible solution is then implemented. The Level 3-controller is written in Python 3. The communication of the optimization results and the sensor measurements between Matlab and Python is facilitated via shared csv files. A Python OPC-UA (Leitner and Mahnke, 2006) client is used for the communication with sensors and actuators of the heat pump and the building. The heat pump takes as an input the relative compressor speed (as discussed above) and uses an internal controller to track it.

<sup>4</sup>The second compressor stage of the heat pump is activated if the relative speed exceeds the limit of 50%.

<sup>5</sup>We conducted preliminary numerical studies to investigate how the solution converges with time.

The heating demand forecast is performed at midnight (for Level 1) and then every 15 minutes (for Level 2) with an ANN and the correction methods presented in Section 3.3 and Chapter 2. The correction based on error-autocorrelation is applied with every new forecast, while the online retraining is done only at midnight. The ANN model uses as inputs the forecast ambient temperature (broadcast by MeteoSwiss and updated every 12 hours), the hour of the day (which is one hot encoded), the measured heating demand one day ago at the same time, the measured heating demand one week ago at the same time, and a binary variable that indicates whether it is a working or a non-working day. The ANN model is implemented in Python 3 with Keras (Chollet, 2018). It is a feed-forward network with two hidden layers and 8 nodes per hidden layer with Rectified Linear Units (ReLU) as activation functions. Just short of three years of historical data (sampled in 15 minute intervals) were used for training using the optimizer *adam* (Kingma and Ba, 2017) with the standard learning rate of 0.001, a batch size of 1, and 10 epochs. This configuration corresponds to the one presented and validated in Chapter 2.

The experiment was conducted on three consecutive days from the 25th of February 2020, 11.45 am, to the 28th of February, 11.45 am. As a regulation signal, the RegD signal by PJM from the 27th of January 2019 was used for all three days.<sup>6</sup>

### 3.4.2 Results

Figures 3.5 - 3.7 show the results of the three-day experiment, split in single days for readability. Note that the time axis is shifted by 11 hours and 45 minutes, to virtually let the experiment start at midnight. In each of the figures, plot (a) depicts the real heating demand of the NEST building in dashed blue, the forecast conducted at midnight for Level 1 in orange, and the forecasts for Level 2, which are conducted every 15 minutes, in transparent grey. The initial forecast (orange) predicts the trend of the heating demand well, confirming the results of Bünning et al. (2020a) and Chapter 2. The correction based on error-autocorrelation is visible, whenever the previous forecast significantly differs from the measured heating demand. This is the case at 00:00 at the beginning of the second day, Figure 3.6, for example. At this point, the initial forecast of 23 kW (orange line) differs from the measured demand of 20 kW after one interval. The next corrected forecast at 00:15 (cyan line) therefore starts at 20 kW. The following corrected forecast at 00:30 (red line) starts at the measured demand of 16 kW. Both corrected forecasts merge back into the initial forecast over the course of the day because the error-autocorrelation also decreases with time. It can also be seen that the experiment covers a range of heating demands from 10 kW, which is below the minimum thermal capacity of the heat pump, to 45 kW, which is close to the maximum capacity of the

---

<sup>6</sup>This choice was made for convenience, as the signal starts and ends with a value of 1 and is thus continuous when repeated.




Figure 3.5: Experimental results of day 1. (a): measured heating demand in dashed blue, daily forecast in orange, 15-minute forecasts in transparent grey, (b): offered reserves in blue, (c): average tank temperature in black, temperature constraints in dotted black, individual layer temperatures in transparent colours, (d): set point for electrical power in orange, measured electrical reserves in dashed blue (no data during hours where no reserves are offered, potential power range due to regulation signal in dotted red), (e): PJM performance score in blue (no data during hours where no reserves are offered), 20-hour moving average of performance score in orange, qualification limit in dotted orange, operation limit in dotted red.




Figure 3.6: Experimental results of day 2. (a): measured heating demand in dashed blue, daily forecast in orange, 15-minute forecasts in transparent grey, specifically mentioned forecasts in cyan and red, (b): offered reserves in blue, (c): average tank temperature in black, temperature constraints in dotted black, individual layer temperatures in transparent colours, (d): set point for electrical power in orange, measured electrical reserves in dashed blue, potential power range due to regulation signal in dotted red, (e): PJM performance score in blue (no data during hours where no reserves are offered), 20-hour moving average of performance score in orange, qualification limit in dotted orange, operation limit in dotted red.




Figure 3.7: Experimental results of day 3. (a): measured heating demand in dashed blue, daily forecast in orange, 15-minute forecasts in transparent grey, (b): offered reserves in blue, (c): average tank temperature in black, temperature constraints in dotted black, individual layer temperatures in transparent colours, (d): set point for electrical power in orange, measured electrical reserves in dashed blue, potential power range due to regulation signal in dotted red, (e): PJM performance score in blue (no data during hours where no reserves are offered), 20-hour moving average of performance score in orange, qualification limit in dotted orange, operation limit in dotted red.

heat pump using only one compressor stage.

Plot (b) in the figures shows the reserves offered during the experiment. The reserves are either zero, or between 0.4 kW and 2.3 kW, which are the lower and upper limits. The upper limit is set as a constraint in the optimization problem, but is also a result of offering symmetric reserves and the electrical capacity of the heat pump being in the range of 8.2 to 12.8 kW. During the first day (Figure 3.5) very little reserve is offered. This is for two reasons. First, the heat pump is frequently switched off because the demand is low. (See also plots (d) in Figure 3.5). In this case, constraint (3.14e) forces the corresponding elements in  $D_w$  and  $D_{\delta,e}$  to be zero, which means that recourse on uncertainties is no longer possible. Second, whenever the heat pump is on, it operates at the lower capacity limits, which for symmetric reserves results in offering no reserves. During the second (Figure 3.6) and the third day (Figure 3.7) of the experiment, reserves are offered during most of the 15 minute intervals. On day 1, 3.1% of the electricity consumed is flexible, on day 2 and 3, 14.9% and 19.1% are flexible respectively. The average of all three days is 13.4%. In addition to being constrained by the thermal inertia of the storage, the reserves are also limited by the region of operation of the heat pump, specifically to 2.3 kW in this particular case. On days one, two and three 4%, 29% and 41% of this potential are exploited respectively. Considering the small size of the buffer storage compared to the thermal demand - on a day with an average heating demand of 35 kW, the thermal capacity of the buffer is just 3% of the daily demand - this is a considerable outcome. As we will see in Section 3.5, it is mostly thanks to the use of affine policies.

Plot (c) shows the average storage temperature in solid black, the temperature constraints at 28 °C and 38 °C in dotted black and the six temperature measurements at different heights of the storage tank in transparent colors. The average temperature stays between the constraints for most of the time, except one 30 minute instance between 7.30 and 8.30 on day two, Figure 3.6. However, during this time the heating demand of the building could still be served, as the upper temperature layer in the storage tank (transparent blue) was above 28 °C at all times, which is a result of the average storage temperature being a lower bound for the temperature of the top water layer in the storage tank.<sup>7</sup> The average temperature stays relatively close to the lower constraint most of the time as a result of the optimization: unless needed for reserves, temperatures above the minimum mean unnecessary consumption of electricity. During the first day (Figure 3.5), where the heating demand is below the minimum thermal capacity of the heat pump, this results in a hysteresis-like behaviour because the heat pump is switched on and off frequently. Between 8.00 and 12.00 on the second day, when the heat pump is switched on due to the reserves offered, the controller regulates the base load  $u^0$  such that the temperature tracks the lower constraint. However, at times where the heating

---

<sup>7</sup>Close investigation of the different temperature layers also shows that swapping of temperature layers indeed does not occur and the assumption for the storage model holds.




Figure 3.8: Results tracking performance, (a): set point for electrical power in orange, measured electrical reserves in dashed blue, potential power range due to regulation signal in dotted red, (b): regulation signal, (c): tracking error, (d): delay score in green, accuracy score in orange, precision score in red, composite score in dotted blue.

demand of the building is substantially overestimated by the initial demand forecast, the average storage temperature rises. This can be seen in the period between 12.00 and 18.00 on the third day (Figure 3.7).

Plot (d) of Figures 3.5 - 3.7 shows the set point for the electrical power of the heat pump in orange, the actual measured power of the heat pump in dashed blue and the possible range of the power due to the offered reserves in dotted red. As the results are difficult to read in this scale, an excerpt of this plot (9.00 to 11.00 of the second day) is shown in the subplot (a) of Figure 3.8. It can be seen that the heat pump delivers the reserves offered by modulating the electricity between the limits defined by the base load and the offered reserves (dotted red) following the regulation signal depicted in Figure 3.8 (b). The effects of different sizes of offered reserves can also be observed. From 9.00 to 9.30, where no reserves are offered, the heat pump does not exactly follow the set point; the resulting tracking error is also evident in subplot (c) of

Figure 3.8. There are two reasons for this tracking error. First, there is measurement noise of approximately  $\pm 200$  W. Second, the heat pump’s internal controller only accepts integer set points for relative compressor speeds (e.g. 34% and 35%, but not 34.5%), which leads to a discontinuous control signal. After 9.30, a range of different reserves is offered, visible from the span between the red dotted lines. Visible from 10.00 to 10.15, large reserves lead to bigger tracking errors because of the ramping limits of the heat pump. Large tracking errors also occur when large steps in the base set-point for the heat pump appear (at 10.15).

Despite these tracking errors, the performance score<sup>8</sup> of the TSO PJM (PJM, 2019) is better when higher amounts of reserves are offered. This is shown in subplot (d) of Figure 3.8. Here, the green line depicts the *delay score* (time delay between reserve signal and system response), the orange line depicts the *accuracy score* (correlation between the reserve signal and the system response), the red line depicts the *precision score* (error between reserve signal and system response) and the dashed blue line depicts the *composite score* (average of the three). The scores are averaged over one hour intervals and are normalized in the interval  $[0, 1]$ , with 1 being the best score. While the delay score is constantly high, both the accuracy score and especially the precision score become worse when the reserves offered are low, because the error relative to the offered reserves becomes large. Especially in the one-hour intervals where only small reserves are offered (or a combination of no reserves and small reserves), the composite performance score becomes low.

With respect to the whole experiment, this result is not problematic, as can be seen by going back to Figures 3.5 - 3.7. In plot (e), the composite performance score is shown in blue. The orange line depicts the 20-hour moving average of the performance score. This is the metric used by PJM to judge whether a device or power plant is suitable for their reserve product. In the qualification phase, the limit for the performance score is 0.75 (dotted orange), while in the operational phase, the limit is lowered to 0.4 (dotted red). It can be seen that the performance score in this experiment is always well above both limits.

## 3.5 Numerical case study

To better understand the performance of the control scheme, we have further analyzed it in two numerical experiments.

---

<sup>8</sup>We note that the performance score is calculated with a Macro Excel sheet provided by PJM and could not be replicated with our own calculations. We refer the interested reader to the original source (PJM, 2019).



### 3.5.1 Use of affine policies compared to open-loop MPC

In the first numerical experiment, we compare the solution of Level 1 with affine policies (as presented in Section 3.3) to a Level 1 scheme based on standard open-loop MPC without feedback policies for various constant heating demands  $v$ . The optimization problem for Level 1 for open-loop MPC is

$$\min_{\substack{x, r, u^0, u_{\text{th}}, \\ z, \tilde{z}, \epsilon}} f^{\text{el}\top} u^0 - f^r\top r + \lambda\top \epsilon \quad (3.15\text{a})$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (3.15\text{b})$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + \bar{w} \odot r), \quad (3.15\text{c})$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (3.15\text{d})$$

$$zU_{\min} \leq u^0 + w \odot r \leq zU_{\max}, \quad (3.15\text{e})$$

$$\tilde{z}R_{\min} \leq r \leq \tilde{z}R_{\max}, \quad (3.15\text{f})$$

$$z, \tilde{z} \in \mathbb{Z}_2^N, \quad (3.15\text{g})$$

$$\epsilon \geq 0, \quad (3.15\text{h})$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta \in \Delta, \forall e \in E. \quad (3.15\text{i})$$

The comparison is made for constant heating demands between 5 kW and 50 kW, in steps of 5 kW. All parameters for the optimization schemes are the same as in Section 3.4. Figure 3.9 and Figure 3.10 show the reserves offered and the cost functions (including the  $\lambda$  term) for both approaches. It can be seen in Figure 3.9, that without using feedback policies, the reserves offered are close to zero for most demands. This is due to the build-up of uncertainty in the state  $x$  (storage temperature) over the horizon of the optimization problem in combination with a relatively small buffer, compared to the magnitude of the heating demand. In contrast, Level 1 with affine policies is able to offer reserves in most cases except when approaching the upper and lower capacity limits of the heat pump. From Figure 3.10 it can be seen that the value of the cost function is significantly lower when affine policies are used for all heating demands below 45 kW. Above 45 kW, there is no difference because the heat pump will always work at maximum capacity. We note however, that the played-out costs (with the MPC re-optimizing every 15 minutes) would have different results for both cases, which significantly depend on the uncertainty realizations.<sup>9</sup>

---

<sup>9</sup>While the played-out behaviour would certainly be an interesting result to study, the computational effort is not feasible.

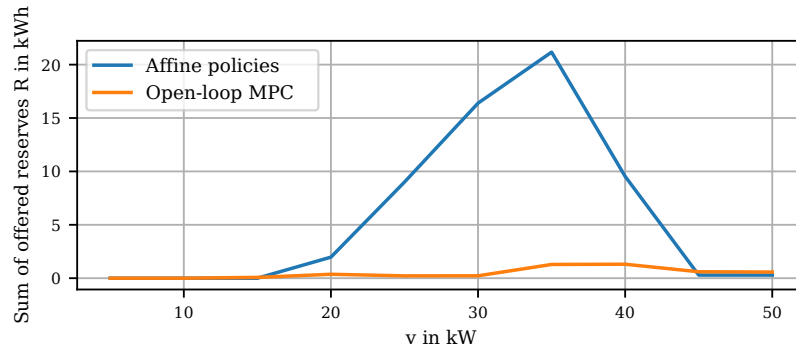


Figure 3.9: Reserves offered in Level 1: affine policies (blue) vs. open-loop MPC (orange).

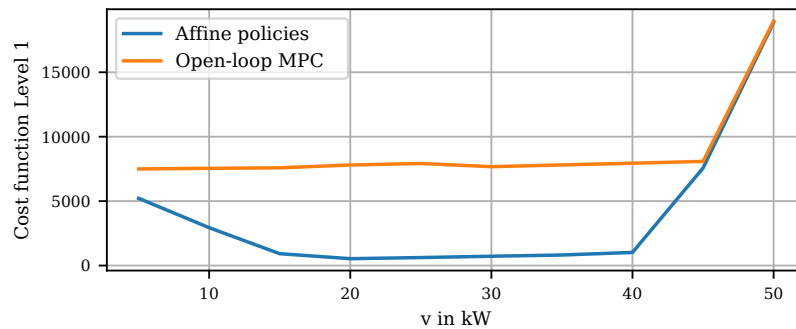


Figure 3.10: Cost in Level 1: affine policies (blue) vs. open-loop MPC (orange).

### 3.5.2 Use of affine policies compared to a system with perfect knowledge

In a second numerical experiment, we compare the performance of Level 1 using affine policies with an omniscient Level 1 solution that has perfect knowledge of all uncertainty realizations at the time of optimization. In this case, the optimization problem becomes

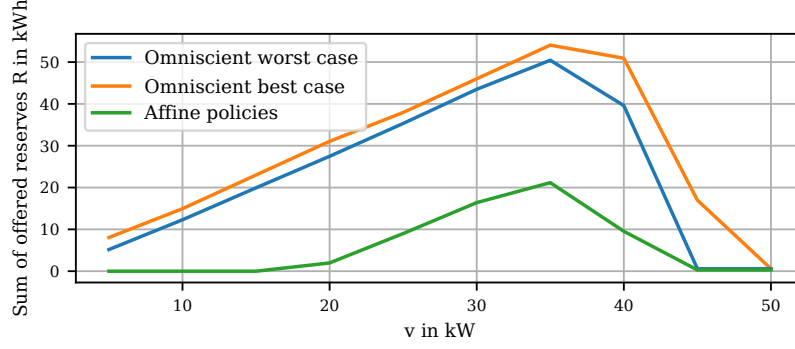


Figure 3.11: Reserves offered in Level 1: affine policies (green) vs. perfect knowledge (best case in orange, worst case in blue).

$$\min_{\substack{x, r, u^0, u_{\text{th}}, \\ z, \tilde{z}, \epsilon}} f^{\text{el}\top} u^0 - f^r \top r + \lambda \top \epsilon \quad (3.16\text{a})$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + (\delta + e)), \quad (3.16\text{b})$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + \bar{w} \odot r), \quad (3.16\text{c})$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (3.16\text{d})$$

$$zU_{\min} \leq u^0 + w_{\min} \odot r, \quad (3.16\text{e})$$

$$u^0 + w_{\max} \odot r \leq zU_{\max}, \quad (3.16\text{f})$$

$$\tilde{z}R_{\min} \leq r \leq \tilde{z}R_{\max}, \quad (3.16\text{g})$$

$$z, \tilde{z} \in \mathbb{Z}_2^N, \quad (3.16\text{h})$$

$$\epsilon \geq 0, \quad (3.16\text{i})$$

where  $(\delta + e)$  is drawn from a uniform distribution with the same limits as the uncertainty set  $E \oplus \Delta$ , and  $w_{\min}$ ,  $w_{\max}$  and  $\bar{w}$  are extracted from the regulation signal used (PJM RegD of 27th of January 2019). Here,  $w_{\min}$  and  $w_{\max}$  are the minimum and maximum value of the regulation signal that occurs during a 15 minute interval respectively;  $\bar{w}$  is the average of the interval.

The experiment is conducted for a constant heating demand  $v$  between 5 kW and 50 kW, in steps of 5 kW, with 10000 uncertainty realizations for each  $v$ . Optimization problem (3.16) is solved much faster than the robust counterpart of problem (3.13), requiring less than 5 seconds for convergence.

Figure 3.11 shows the reserves offered for both cases for all different  $v$ . Here, the orange line denotes the best result achieved with the omiscient system with 10000 uncer-

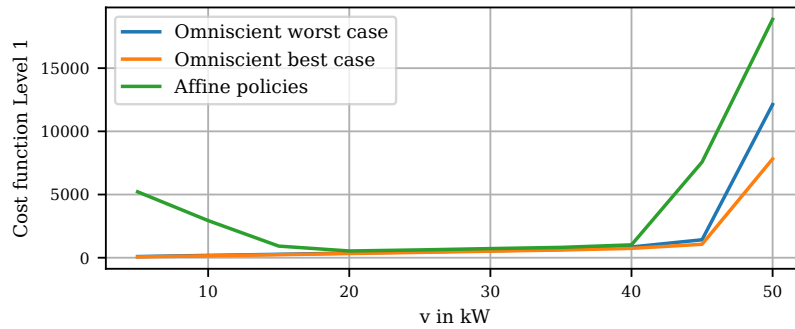


Figure 3.12: Cost in Level 1: affine policies (green) vs. perfect knowledge (best case in orange, worst case in blue).

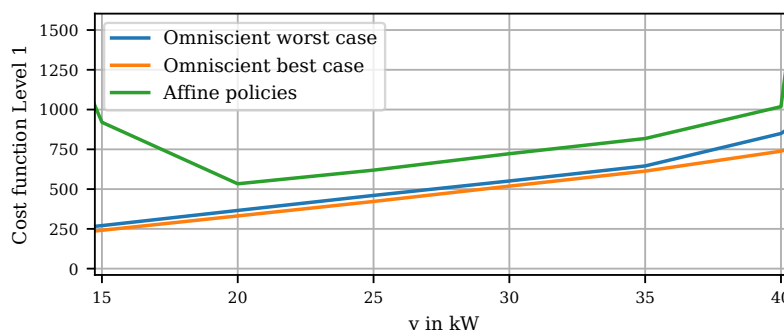


Figure 3.13: Cost in Level 1: affine policies (green) vs. perfect knowledge best case in orange, worst case in blue), in detail.

tainty realizations, while the blue line denotes the worst result.<sup>10</sup> The green line shows the solution of Level 1 with affine policies. It can be seen that for all cases except  $v = 45$  and  $50$  kW, the solutions with perfect knowledge of the uncertainties, offer significantly more reserves than the scheme with affine policies. The results for the cost function, depicted in Figure 3.12 for all  $v$  and depicted in Figure 3.13 for  $v$  between  $15$  and  $40$  kW, show that the scheme with perfect knowledge performs significantly better at very high and very low heating demands, which is due to avoiding the use of the slack variable  $\epsilon$ . Also at intermediate demands, there is an offset between the cost functions of both schemes. This can be explained by the fact that the scheme with perfect knowledge is able to operate right at the storage temperature constraints (maximizing offered reserves or minimizing the base load), while the scheme with affine policies needs to stay at least  $B(\Delta \oplus E)$  away from the storage temperature constraint, even when no reserves are offered. This is the case because affine policies can only react to uncertainties in the timestep after their realization, unlike a system with perfect knowledge, which can plan ahead and can compensate for uncertainties before they are revealed. The larger optimality gap at  $15$  kW can again be explained by the fact that affine policies can only work when  $z$  is non-zero, which is often not the case for low heating demands.

### 3.6 Limitations and future directions

The numerical results show that using affine policies on uncertainties significantly increases reserve provision and decreases cost when compared to using open-loop MPC. Nevertheless, the comparison with the omniscient optimization indicates that there is still room for improvement. One possibility could be the use of more sophisticated policies (for example deflected linear (Chen et al., 2008), or piecewise linear (Georghiou et al., 2015) policies). However, due to causality constraints, no policy can match the performance of an omniscient controller that knows the future.

Moreover, the presented method of only exploiting the thermal inertia of buffer storage should be extended to and compared to a scheme where the inertia of a subset of the connected buildings (willing and able to participate in the reserve provision scheme) is also taken into account, to further investigate the added potential by including building thermal dynamics in the problem formulation. This will be discussed in Chapter 5, after appropriate data-driven models for building thermal dynamics are introduced in Chapter 4.

On a more practical note, a limitation of the approach presented here is that it can only work with variable speed heat pumps, which, although becoming more common,

---

<sup>10</sup>Each set of uncertainty realizations drawn leads to a different cost function and to differing amounts of reserves offered, as each realization either benefits the operation or harms it, compared to the case without uncertainty. We report the best case and worst case among these realizations.

are still relatively rare compared to fixed speed heat pumps. To offer reserves with fixed speed devices, an aggregation mechanism that pools heat pumps of many different buildings would still be necessary in order to follow a continuous reserve signal. A further limitation for air-sourced heat pumps could be the requirement of regular de-freezing cycles. However, as these do not last long and our achieved performance scores are well above the qualification limit, they could be integrated without a change to the control scheme.

### **3.7 Conclusion**

In this chapter, we have combined a three-level control scheme based on robust optimization with affine policies with heating demand forecasting based on ANN and online correction methods to provide frequency regulation reserves with heat pumps and water storage in buildings and district heating settings. The approach works without the necessity of a dynamic building model, which reduces modeling effort compared to including reserve provision in MPC building temperature control. Moreover, it alleviates privacy concerns and has reduced hardware requirements, as no indoor temperature or occupancy measurements are necessary to guarantee indoor comfort.

The experiments on the heat pump and water storage in the NEST demonstrator building have shown that the three-level control approach with affine policies on uncertain variables presented here is viable. The method allows the offering of frequency regulation reserves with a single variable speed heat pump and a small (compared to the overall heating demand) buffer storage. On average, 13.4% of the consumed electricity is flexible as a result of the reserves offered. The performance of tracking the regulation signal is more than sufficient. Our experiments indicate that large buildings or district heating systems equipped with variable speed heat pumps can in principle directly be used for ancillary services, without the need for aggregation, and that the control scheme is robust in practical applications. The heating demand forecasting approach with ANN and online correction methods gives predictions with high accuracy, such that the demand of the building can always be met. The numerical case studies show that, while there is still a performance gap compared to an omniscient controller, the use of affine policies significantly reduces costs and improves reserve provisions, compared to open-loop MPC. In Chapter 5, the method will be extended to a problem formulation that also exploits the thermal inertia of connected buildings.

# Data-driven predictive building control

Models of building thermal dynamics can be used in Model Predictive Control (MPC) schemes to reduce building energy consumption, or to extend electrical reserve schemes like the one presented in Chapter 3, to exploit their thermal inertia for reserves. Because physics-based building models are difficult to obtain as each building is individual, there is an increasing interest in generating models suitable for building MPC directly from measurement data. Machine Learning methods have been widely applied to this problem and validated mostly in simulation; there are, however, few studies on a direct comparison of different models or validation in real buildings in the literature. Methods that are indeed validated in application often lead to computationally complex non-convex optimization problems. Here, we introduce physics-informed Autoregressive–Moving-Average with Exogenous Inputs (ARMAX) models and Machine Learning models based on Random Forests and Input Convex Neural Networks, and compare the resulting convex MPC schemes in experiments on a practical building application, and in a numerical case study. We demonstrate that Predictive Control in general leads to savings between 26% and 49% of heating and cooling energy, compared to the standard hysteresis controller used in the building. Moreover, we show that all model types lead to satisfactory control performance in terms of constraint satisfaction and energy reduction. However, we also see that the physics-informed ARMAX models have a lower computational burden, and a superior sample efficiency compared to the Machine Learning based models. Moreover, even if abundant training data is available, the ARMAX models have a significantly lower prediction error than the Machine Learning models, which indicates that the encoded physics-based prior of the former cannot independently be found by the latter.

## 4.1 Introduction

MPC in buildings has been applied successfully in simulation (Oldewurtel et al., 2012; Touretzky and Baldea, 2014; Chen et al., 2013; Mai and Chung, 2015) and real buildings (Sturzenegger et al., 2016; Hilliard et al., 2017; Castilla et al., 2014; Ma et al., 2014;

(Žáčková et al., 2014; Hammer and Gersmann, 2003) many times with significant reduction of energy consumption compared to the baseline controllers. It can also be used to provide electrical reserves with buildings by exploiting their thermal inertia, as demonstrated in (Vrettos et al., 2018b), which is a second application that is highly relevant to the scope of this thesis. However, models in building MPC are conventionally based on physics (Picard et al., 2015; Sturzenegger et al., 2014), which means that the model is built using principles of heat transfer and thermodynamics, or based on expensive excitation experiments, or a combination of both. Another approach are grey-box methods (Dimitriou et al., 2015; De Coninck et al., 2015; Li et al., 2021), where the basic physical and architectural structure of the considered building is modeled by hand, which also requires a considerable amount of manual work, and parameters are fitted to measurement data. Several recent (Maddalena et al., 2020; Kathirgamanathan et al., 2021; Drgoňa et al., 2020a; Péan et al., 2019; Lee and Karava, 2020; Mariano-Hernández et al., 2021; Gholamzadehmir et al., 2020; Zong et al., 2019) and less recent (Rockett and Hathway, 2017; Henze, 2013; Afram and Janabi-Sharifi, 2014; Serale et al., 2018) reviews provide an excellent overview on the issues related to the use of MPC in buildings, including physics-based modeling pipelines such as the Building Resistance-Capacitance Modeling (BRCM) toolbox (Sturzenegger et al., 2014) or Modelica-based approaches (Picard et al., 2015). A common observation is that developing and maintaining such models is often considered too expensive to justify investment as each building is individual, an issue that potentially hinders the commercial application of MPC in buildings (Sturzenegger et al., 2016).

As a result, data-driven approaches that rely purely on historical measurement data have emerged. Data-driven methods, which are sometimes referred to as Data Predictive Control (DPC) (Smarra et al., 2018), either use models built from measurement data in an MPC framework (Smarra et al., 2018) or compute optimal inputs directly from past and currently measured data (Coulson et al., 2019; Markovsky et al., 2006). As the distinction between these methods and methods based on excitation experiments is blurry, we will use the common term MPC in the following for all predictive methods. Many Machine Learning (ML) methods, such as Artificial Neural Networks (ANN), Random Forests (RF) and Support Vector Machines (SVM) are universal function approximators (Hornik, 1991; Hammer and Gersmann, 2003), and have proven successful in various technical domains (Silver et al., 2017; Wu et al., 2016; Hershey et al., 2017). As they come with the tempting promise to model any system well as long as the underlying data quality is good, they are also natural candidates for MPC in buildings.

There are several studies on the comparison of such modeling methods in the context of data-driven MPC in simulation. While some authors (Mugnini et al., 2020; Wang et al., 2019; Chen et al., 2019) compare purely data-driven methods such as ANN, RF and SVM to grey-box methods such as resistor-capacitor (RC) models, others (Picard et al., 2016) compare RC methods to white-box models, or subspace identification methods



to ARMAX (Ferkel and Jan Široký, 2010). Other studies apply data-driven MPC for buildings in practice but do not compare different modeling methods. For example, Ferreira et al. (2012) use an ANN-based MPC controller on a set of rooms in a real university building. The controller saves up to 50% of energy compared to the baseline controller. Aswani et al. (2012) use switched linear input-output models, which are trained based on excitation experiment data, to control a university building with an MPC framework. In an eight day experiment, the controller saves a significant amount of energy compared to the baseline controller, while providing similar comfort levels. Finck et al. (2019); Jain et al. (2020) and Yang et al. (2020) use ANN-based MPC controllers with varying types of cost functions on real building applications with varying purposes, such as minimizing energy use or maximizing demand flexibility. The review article of Kathirgamanathan et al. (2021) concludes that the validation and comparison of data-driven MPC controllers in practice is not appropriately addressed in the literature so far.

In the spirit of Mania et al. (2018), it is also unclear how the considered Machine Learning based methods models compare to simpler identification methods, such as ARMAX models identified through linear regression. While ARMAX models are not universal function approximators, they do have the trait of being physically interpretable, which is often not the case for ML methods. There is a general growth in research on physics-informed data-driven methods (see e.g. (Raissi et al., 2017, 2019; Sharma et al., 2018; Manek and Kolter, 2020; Lutter et al., 2019; Márquez-Neila et al., 2017)), but very little work has been done so far in the domain of building energy control. Drgoňa et al. (2020b) introduce physics-constrained RNN to model the thermal dynamics of buildings and use information about the general model structure of buildings to structure the neural dynamics models, constrain the eigenvalues of the model, and use penalty methods to impose physically meaningful boundary conditions to the learned dynamics. The method is applied to an open-loop prediction on a data set of a real 20-zone building. The authors find that the prediction accuracy is significantly improved compared to not-constrained RNN models. The method is not tested in closed-loop MPC. To the best of our knowledge this is the only published study on the topic of physics-informed data-driven models in the building control domain. There is no study available on a closed-loop implementation on a physical building.

Besides the addressed lack of physics-informed data-driven models, the body of studied literature has shown that there is a general lack of comparisons of the performance of different data-driven models in MPC for practical building applications (Kathirgamanathan et al., 2021). Moreover, in cases where data-driven methods are indeed tested on real systems, the resulting MPC problem is often non-convex and has to be solved with non-linear solvers. While this approach might be suitable for single building energy management, the resulting computational complexity of solving the MPC problem online likely leads to intractable problems if more complex system architectures or ap-

plications are considered. Examples for such problems are coordinated building control (Yang et al., 2020), or electric reserve provision as discussed in Chapters 3 and 5.

### 4.1.1 Contribution

In this chapter, we therefore focus on the development and practical validation of data-driven methods that lead to convex MPC problem formulations. We introduce two Machine Learning based data-driven models for building thermal dynamics. The first one is a combination of Random Forests (RF) and linear regression, based on the work of (Smarra et al., 2018). The second one is based on Input Convex Neural Networks (ICNN), which were extended from (Amos et al., 2017) to ensure convexity in the face of recursive evaluation of the networks. Under some assumptions on the cost function and state constraints, both RF and ICNN models lead to convex optimization problems that can be solved to global optimum in real time. Moreover, we introduce physics-informed ARMAX models by exploiting the physical interpretability of ARMAX models and their relation to the physical dynamics of buildings.

We conduct a series of heating and cooling experiments on an occupied apartment (see lit apartment on second floor in Figure 2.2, Chapter 2), with the different models embedded in an MPC framework, and compare them qualitatively in terms of constraint satisfaction and exploitation, and in terms of the computational burden, i.e the online optimization time and memory requirements. Moreover, we quantitatively evaluate the energy savings achieved with MPC compared to the baseline hysteresis controller used in the apartment, which amount to 26% to 49%. In a numerical experiment on the basis of historical measurement data from the apartment, we also compare the models in terms of sample efficiency, i.e. how much data a good model needs for training, and multi-step prediction accuracy. While all models perform well on the task of predictive control in the application case study and outperform the baseline hysteresis controller in terms of energy consumption, the physics-informed ARMAX models show a considerably lower computational burden, better sample efficiency and better prediction accuracy in the numerical evaluation. On one hand, our results show that any model with reasonable accuracy will be suitable for predictive control. On the other hand, they demonstrate that ML-based models do not have any advantage over a physics-informed ARMAX model in the given case, i.e. they do not identify potential non-linearities such as time related solar gains or occupancy effects. On the contrary, the stronger physical prior of the ARMAX models, besides leading to higher sample efficiency, is not achieved by the ML models, even if abundant training data is available for training.

### 4.1.2 Structure

The remainder of this chapter is structured as follows. In Section 4.2, we discuss the concept of MPC with input-output models. In Section 4.3, we introduce the models based on RF, ICNN and physics-informed ARMAX. Section 4.4 describes the apartment that we use as a test bed and its heating and cooling system. The results of the experiments of all models applied to the test case are presented and discussed in Section 4.5, while in Section 4.6, we investigate the sample efficiency of the models. Section 4.7 summarises our findings and provides directions for future work.

## 4.2 Problem statement

MPC is a control scheme where a constrained optimization problem is solved repeatedly to find optimal control inputs over a receding horizon. Besides state-space formulations, the problem can be formulated with an input-output model, where previous outputs, inputs and disturbances are measured, and the optimization problem

$$\min_{u,y} \sum_{k=0}^{N-1} J_k(y_{k+1}, u_k) \quad (4.1a)$$

$$\text{s.t.} \quad y_{k+1} = f(y_{k-}, u_{k-}, d_{k-}) \quad (4.1b)$$

$$(y_{k+1}, u_k) \in (\mathcal{Y}_{k+1}, \mathcal{U}_k) \quad (4.1c)$$

$$\forall k \in [0, \dots, N-1],$$

is solved at discrete time instants. Here, the variables  $y$ ,  $u$  and  $d$  denote the system outputs, control inputs and disturbances respectively. The variable  $k$  denotes the time step in the horizon  $N$ ,  $J_k$  is the stage cost, and  $f$  denotes the model describing the system dynamics as a function of autoregressive terms of outputs, and moving average terms of inputs and disturbances (or their forecasts), denoted by the subscript  $-$ . Output and input constraints are formulated with the sets  $\mathcal{Y}_k$  and  $\mathcal{U}_k$ .

Problem (4.1) is solved every time a new measurement of the output is available; after each optimization, the first element of the optimal input sequence,  $u_0^*$ , is applied to the system and the process is repeated. Input-output models are suitable for building control, as there are usually no constraints on hidden system states (for example wall temperatures).

Problem (4.1) describes a basic MPC scheme, but many alternatives have been explored in the literature, for example applying multiple steps of the optimal sequence before repeating the optimization, treating uncertainty in the disturbance forecast in a worst-case or stochastic way, using state space formulations for the system dynamics and

state estimators during controller application, etc. The interested reader is referred to (Kouvaritakis and Cannon, 2016) for information on many of these alternatives.

The control performance generally improves with solving (4.1) more frequently (i.e. the controller has a smaller sampling time) and a larger prediction horizon  $N$ . This, however, implies that less time is available to solve a more complex optimization problem. As convex optimization problems can usually be solved more efficiently, formulations where the cost and constraints in (4.1) are convex with respect to the decision variables are often preferable. Indeed, the art of MPC design is making design choices to master this trade-off between computation and control performance.

## 4.3 Methodology

In this section, we describe our models based on RF, ICNN, and physics-informed AR-MAX models.

### 4.3.1 Random Forests with linear regression leaves

The presented model, extended from the work of Smarra et al. (2018), is an input-output model  $y = f(X_c, X_d)$ , where  $X_c$  are the controllable inputs, i.e. control inputs  $u$  (such as valve positions, supply temperatures, etc.) in the prediction horizon, but not those that lie in the past (as they cannot be decided on any more). The vector  $X_d$  denotes all uncontrollable model inputs, i.e. disturbances (for example ambient temperature, solar irradiation, time features, etc.), previously measured room temperatures and control inputs  $u$  that lie in the past. The model is built in two steps as follows.

#### Fitting of random forests

First, a random forest  $g(X_d)$  is built for each prediction step in (4.1), i.e.  $N$  forests, which map the uncontrollable inputs  $X_d$  to a finite set of leaves  $1, \dots, \mathcal{L}$ . Random forests are ensembles of regression trees - a form of decision tree to approximate continuous functions. A decision tree is built by splitting the training set  $\mathcal{X}_d$  of the input  $X_d$  into partitions  $p$ , yielding  $\mathcal{X}_d^p$ , in each of which the output  $y$  is approximated with a constant value  $\bar{y}^p$  (which is the mean of all measured outputs  $y \in \mathcal{Y}_d^p$  corresponding to  $\mathcal{X}_d^p$ ). This constant value at the bottom of a tree is called a leaf. The splitting variable and splitting point defining the partitions  $\mathcal{X}_d^p$  are obtained with the help of a greedy heuristic:

Consider a splitting variable  $v \in \mathcal{X}_d^p$  and a splitting point  $s$ , the pair of half planes  $R$  defined by these are

$$R_1(v, s) = \{X_d \in \mathcal{X}_d^p | v \leq s\} \text{ and } R_2(v, s) = \{X_d \in \mathcal{X}_d^p | v > s\}. \quad (4.2)$$

We seek the splitting variable  $v$  and splitting point  $s$  that solve

$$\min_{v,s} \left( \min_{\bar{y}^{R1}} \sum_{X_{d,j} \in R_1(v,s)} w_j (y_j(X_{d,j}) - \bar{y}^{R1})^2 + \min_{\bar{y}^{R2}} \sum_{X_{d,j} \in R_2(v,s)} w_j (y_j(X_{d,j}) - \bar{y}^{R2})^2 \right) \quad (4.3)$$

by iterating over  $v$  and  $s$ . Here,  $y_j(X_{d,j})$  denotes the measured output corresponding to a specific input  $X_{d,j}$ . The variable  $w_j$  can be used to apply sample-weighting techniques. If no sample weighting is applied, its value is 1 for all inputs.

A random forest of  $\mathcal{T}$  trees is built by bootstrapping (Efron, 1979)  $\mathcal{T}$  subsets of samples from the available training data and building a tree based on each subset. The function evaluation of the forest for  $X_d$  is then done by evaluating each tree, giving rise to  $\mathcal{T}$  different  $\bar{y}(X_d)$  which are then averaged.  $\mathcal{T}$  is independent of  $N$  and needs to be chosen heuristically. For more details on the training process please refer to the original source (Smarra et al., 2018) and to general literature on random forests (Louppe, 2014).

### Linear regression leaves

Second, instead of using the constant values  $\bar{y}(X_d)$  as predictors, in each leaf of the forest, a function  $h_i(X_c)$  with  $i \in 1, \dots, \mathcal{L}$  is defined on the basis of the controllable inputs, which maps  $X_c$  to  $y$ . Smarra et al. (2018) fit an affine model to approximate  $y$  with  $X_c$  (to relate the response of the room temperature to the control inputs) with least squares in each leaf  $i$ . For the prediction of the future temperature of a single room  $y$  at time  $k + 1$ , this gives

$$y_{k+1} = \beta_{k,i,X_d} + \sum_{n=0}^k \beta_{k,i,X_c,n} u_n + e, \quad (4.4)$$

in which  $\beta_{k,i,X_d}$  and  $\beta_{k,i,X_c,n}$  denote the fitted coefficients and  $e$  the model error. The temperature is therefore an affine function of all control inputs from the time instant when the prediction is made, to the predicted time instant  $k + 1$ . As each leaf  $i$  is based on a different partition  $\mathcal{X}_d^p$ ,  $\beta_{k,i,X_d}$  and  $\beta_{k,i,X_c,n}$  are also different for each leaf, thus implicitly dependent on  $X_d$ .

Preliminary results suggested that the proposed approach leads to weak prediction performance for more than three prediction steps. This can be explained by the high dimensionality of the fitting process for larger prediction steps in combination with a lim-

ited amount of training samples per leaf. We therefore model the predicted temperature in the  $k + 1^{th}$  step as

$$y_{k+1} = \beta_{k,i,X_d} + \beta_{k,i,X_c} \sum_{n=0}^k u_n + e \quad (4.5)$$

in which only two coefficients have to be fitted for any number of prediction steps. This approach is less realistic than eq. (4.4), as an early control input will have the same effect on the predicted temperature as a late control input (close to  $k + 1$ ), whereas one would expect the effect of earlier time steps to be weaker. However, preliminary experiments suggested that the advantage of better model fits outweighs the disadvantage of a less realistic model. As we use random forests and not the individual regression trees for predicting  $y$ , the average of all relevant  $\beta_{k,i,X_d}$  and  $\beta_{k,i,X_c}$  is taken when the model is applied in (4.1).

The resulting prediction function is  $f(X_c, X_d) = h_{g(X_d)}(X_c)$ . When the model is applied in MPC as (4.1b), it is not applied recursively. Instead, for each time step  $k$  in the horizon  $N$ , a separate model  $f_k$  is built. Here, only previously measured outputs, but not previously predicted outputs are used as model inputs to keep the optimization problem convex. For example, we could define  $y_1 = f_1(X_c = (u_0), X_d = (y_0, y_{-1}, d_0, d_{-1}))$  for the first prediction, but  $y_2 = f_2(X_c = (u_1, u_0), X_d = (y_0, d_1, d_0))$  for the second prediction step; i.e.  $f_2$  is not a function of  $y_1$ , only of  $u_1$ . This is done to maintain the convexity of the MPC controller. The optimization problem is solved by looking up the leaves in the forest on the basis of measurements and forecasts of  $X_d$  and collecting the appropriate functions  $h_i(X_c)$  (i.e eq. (4.5)) first, and second by solving the resulting optimization problem.

### 4.3.2 Input convex neural networks

We extend the work of Amos et al. (2017), to describe building dynamics with a Neural Network that allows a convex formulation of problem (4.1) in the absence of lower output constraints. The model is an input-output model  $y = f(X_{cvx}, X_{ncvx})$ , where the output is convex with respect to  $X_{cvx}$  (for example decision variables such as room temperatures and control inputs), but not necessarily with respect to  $X_{ncvx}$  (for example disturbances such as solar irradiation or ambient temperature), even for multi-step predictions. The derivation of the model is explained in the following.

#### Convex one-step ahead prediction

Amos et al. (2017) introduced an architecture for feed-forward Neural Networks where the scalar output of the network is convex with respect to all inputs (Fully Input Convex

Neural Network or FICNN), or with respect to a subset of the inputs (Partially Input Convex Neural Network or PICNN). These networks are therefore promising candidates for  $f$  in problem (4.1).




Figure 4.1: Schematic of a Fully Input Convex Neural Network

Figure 4.1 shows the schematic of an FICNN. It shows a  $L$ -layer fully connected network in which the output of each layer follows

$$z_{i+1} = g_i(W_i^{(z)} z_i + W_i^{(x)} X_{\text{cvx}} + b_i), \quad (4.6)$$

with  $z_0 = 0$  and  $W_0^{(z)} = 0$ , where  $g_i$  denotes the activation function for layer  $i \in [1, \dots, L-1]$ ,  $W_i^{(z)}$  and  $W_i^{(x)}$  denote the network's weights, and  $b_i$  denotes a constant bias. Both weights and biases are model parameters that are determined during network training. The network output of the last layer  $z_L$ , which is also the model output  $y$ , is convex with respect to the elements of the input vector  $X_{\text{cvx}}$  if all  $W_{1:L-1}^{(z)}$  are non-negative and all activation functions  $g_i$  are convex and non-decreasing (Amos et al., 2017).




Figure 4.2: Schematic of a Partially Input Convex Neural Network

Similarly, Figure 4.2 depicts a PICNN. Here, the output of each network layer follows

$$v_{i+1} = \tilde{g}_i(\tilde{W}_i v_i + \tilde{b}_i),$$

$$z_{i+1} = g_i \left[ W_i^{(z)} \left\{ z_i \circ [W_i^{(zv)} v_i + b_i^{(z)}] \right\} + W_i^{(x)} \left\{ X_{\text{cvx}} \circ [W_i^{(xv)} v_i + b_i^{(x)}] \right\} + W_i^{(v)} v_i + b_i \right],$$

(4.7)

with  $z_0 = 0$ ,  $W_0^{(z)} = 0$  and  $v_0 = X_{\text{ncvx}}$ , where  $\tilde{g}_i$  and  $g_i$  are activation functions,  $\tilde{W}_i$ ,  $W_i^{(z)}$ ,  $W_i^{(zv)}$ ,  $W_i^{(xv)}$ ,  $W_i^{(v)}$  are input weights,  $\tilde{b}_i$ ,  $b_i$ ,  $b_i^{(z)}$ ,  $b_i^{(x)}$  are constant biases, and  $\circ$  denotes the Hadamard product. Under the condition that all weights  $W_{1:L-1}^{(z)}$  are non-negative and all activation functions  $g_i$  are convex and non-decreasing, the model output  $y$ , i.e. the output of the last layer  $z_L$ , is convex with respect to the elements of the input vector  $X_{\text{cvx}}$  (but not necessarily with respect to the elements of  $X_{\text{ncvx}}$ ) (Amos et al., 2017).

FICNN and PICNN guarantee input-output convex behaviour, making them promising candidates for MPC schemes, as mentioned above. However, when ICNNs are applied for  $f$  in problem (4.1), for example as  $y_{k+1} = f(X_{\text{cvx}} = (y_k, u_k, d_k))$ , it becomes evident that there is an issue with convexity in multi-step ahead prediction. While the first step  $y_1 = f(y_0, u_0, d_0)$  is a convex function of the decision variables, the second one is not guaranteed to be. For example, if a 1-layer FICNN is considered, the output at timestep 2 follows

$$y_2 = f(y_1, u_1, d_1) = g_0(W_0^{(x)}(y_1, u_1, d_1) + b_0). \quad (4.8)$$

As the elements of  $y_1$  are convex functions of the elements of  $y_0$  and  $u_0$ , take  $y_1 = y_0^2 + u_0^2$  as a scalar example, and  $W_0^{(x)}$  can attain any value (thus also a negative one), the term  $W_0^{(y)}(y_1, u_1, d_1)$  can become a concave function in the elements of  $y_0$  and  $u_0$ , for example  $(-1)(y_0^2 + u_0^2)$ . The output  $y_2 = f(y_1, u_1, d_1)$  is thus not guaranteed to be convex with respect to the elements of  $y_0$  and  $u_0$ .

### ICNN for multi-step ahead prediction

To make ANNs input convex in the face of multi-step ahead prediction, Chen et al. (2019) have presented a solution for Fully Input Convex Recurrent Neural Networks. Here, we present a solution for Fully and Partially Input Convex Feed-Forward Neural Networks. We accomplish this through additional constraints in the network architecture and the use of ReLu (Nair and Hinton, 2010) activation functions in appropriate places in the network topology. Our approach readily extends to FICNN, where it becomes the feed-forward counterpart of the approach developed by Chen et al. (2019) for recurrent neural networks.

**Proposition 1:** *Consider the two FICNNs  $f_1(X_{\text{cvx},1})$  and  $f_2(X_{\text{cvx},2})$  as defined in eq. (4.6). The composition  $f_2(X_{\text{cvx},2} = (\hat{X}_{\text{cvx},2}, f_1(X_{\text{cvx},1})))$  is a convex function with respect to the elements of the vector  $X_{\text{cvx},1}$ , if all weights  $W_i^{(z)}$  and  $W_i^{(x)}$  are non-negative and all functions  $g_i$  are convex and non-decreasing. Here,  $\hat{X}_{\text{cvx},2}$  are the inputs unrelated to  $f_1$ .*



The proof follows from the fact that non-negative sums of convex functions are convex and that compositions of a convex function and a convex non-decreasing function are also convex. Looking at eq. (4.6), the elements of  $W_i^{(z)} z_i$  are convex functions assuming that  $z_i$  is convex in its inputs and all elements of  $W_i^{(z)}$  are non-negative. The parameter  $b_i$  is a constant. Generally, the term  $W_i^{(x)} X_{\text{cvx}}$  is convex in  $X_{\text{cvx}}$  for any  $W_i^{(x)}$  if  $X_{\text{cvx}}$  is constant because a linear function with negative gradient is also convex (Amos et al. (2017)). Here, we require the elements of  $W_i^{(x)}$  to be non-negative, because  $X_{\text{cvx}}$  itself might be a convex function in the form of an ICNN (from a previous prediction step). The term  $(W_i^{(z)} z_i + W_i^{(x)} X_{\text{cvx}} + b_i)$  is therefore a non-negative sum of convex functions and  $g_i(W_i^{(z)} z_i + W_i^{(x)} X_{\text{cvx}} + b_i)$  a composition of a convex function and a convex non-decreasing function (for example the commonly used ReLU function  $g_i = \max(x, 0)$ ). For the example of eq. (4.8), as  $W_0^{(x)}$  is now constrained to be non-negative,  $W_0^{(x)}(y_1 = f(y_0, u_0, d_0))$  is convex in the elements of  $y_0$  and  $u_0$ , if  $y_1$  is convex in the elements of  $y_0$  and  $u_0$  (which is the case). Thus,  $y_2$  is also convex in the elements of  $y_0$  and  $u_0$ . The network output  $z_L$  is a convex non-decreasing function of the elements of the input  $X_{\text{cvx}}$ . An alternative proof can be constructed from showing that a FICNN is a convex non-decreasing function and using the composition rule of convex and convex non-decreasing functions.

In the case of PICNN, we propose the following structure for the outputs of the network layers for input convex multi-step prediction,

$$\begin{aligned} v_{i+1} &= \tilde{g}_i(\tilde{W}_i v_i + \tilde{b}_i), \\ z_{i+1} &= g_i \left[ W_i^{(z)} \left( z_i \circ g_i^{(zv)} [W_i^{(zv)} v_i + b_i^{(z)}] \right) + W_i^{(x)} \left( X_{\text{cvx}} \circ g_i^{(xv)} [W_i^{(xv)} v_i + b_i^{(x)}] \right) \right. \\ &\quad \left. + W_i^{(v)} v_i + b_i \right], \end{aligned} \quad (4.9)$$

where the activation functions  $g_i^{(zv)}$  and  $g_i^{(xv)}$  are added compared to eq. (4.7).

**Proposition 2:** *Consider the two PICNNs  $f_1(X_{\text{ncvx},1}, X_{\text{cvx},1})$  and  $f_2(X_{\text{ncvx},2}, X_{\text{cvx},2})$  as defined in eq. (4.9). The composition  $f_2(X_{\text{ncvx},2}, X_{\text{cvx},2} = (\hat{X}_{\text{cvx},2}, f_1(X_{\text{ncvx},1}, X_{\text{cvx},1})))$  is a convex function with respect to the elements of  $X_{\text{cvx},1}$  (and  $X_{\text{cvx},2}$ ), but not necessarily with respect to the elements of  $X_{\text{ncvx},1}$  (and  $X_{\text{ncvx},2}$ ), if all weights  $W_i^{(z)}$  and  $W_i^{(x)}$  are non-negative, all functions  $g_i^{(zv)}$  and  $g_i^{(xv)}$  map to a non-negative value and the function  $g_i$  is convex and non-decreasing. Here,  $\hat{X}_{\text{cvx},2}$  are the inputs unrelated to  $f_1$ .*

The proof again follows from only applying operations that maintain convexity. Going through eq. (4.9) term by term,  $z_i \circ g_i^{(zv)} [W_i^{(zv)} v_i + b_i^{(z)}]$  is convex in the elements of the inputs of  $z_i$  if  $z_i$  is a convex function because  $g_i^{(zv)}$  maps all negative values of  $[W_i^{(zv)} v_i + b_i^{(z)}]$  to zero. As  $W_i^{(z)}$  is non-negative,  $W_i^{(z)}(z_i \circ g_i^{(zv)} [W_i^{(zv)} v_i + b_i^{(z)}])$  is also convex.

The same argument can be made for the next term  $W_i^{(x)}(X_{\text{cvx}} \circ g_i^{(xv)}[W_i^{(xv)}v_i + b_i^{(x)}])$  regarding the convexity in the elements of  $X_{\text{cvx}}$  and its inputs. Here, we need  $g_i^{(xv)}$  to map all negative values to zero, because  $X_{\text{cvx}}$  is not a constant (as in Amos et al. (2017)), but a convex function itself. As  $g_i$  is convex non-decreasing, the composition of  $g_i$  and before-mentioned terms is a convex function. Eq. (4.9) is thus convex with respect to the elements of the input vectors of any convex functions  $X_{\text{cvx}}$  and  $z_i$ . We note in passing, that other formulations of FICNN and PICNN can be thought of and will give the same result as long as it is ensured that  $z_{i+1}$  is convex non-decreasing in the elements of  $X_{\text{cvx}}$ .

When PICNN or FICNN are applied in the MPC problem (4.1), constraints (4.1b) and (4.1c) define a convex set as long as the outputs are box-constrained and no lower bounds are imposed. This is because the equality constraint can be eliminated. For example, the dynamics  $y_2 = f(y_1, u_1, d_1)$  and a box-constrained output  $y_2 \leq y_{\text{max}}$  can be replaced with  $f(f(y_0, u_0, d_0), u_1, d_1) \leq y_{\text{max}}$ , which defines a convex set because  $f$  is convex non-decreasing. We note, that in many practical cases, the problem could remain convex also in the presence of lower output constraints due to the monotonicity of the dynamics.

### 4.3.3 Physics-informed ARMAX models

The models presented in Subsections 4.3.1 and 4.3.2 fall into the family of Machine Learning methods. In the case of modeling building dynamics, it is tempting to use them to learn the building's behaviour, as some dynamics-related effects are difficult to model from first principles for individual buildings. An example is the heating gain through windows as a function of time, global solar irradiation, window size, and window orientation. However, the ML methods usually do not encode constraints coming from building physics, for example the law that heat flows from warm to cold, the first law of thermodynamics, etc. In the following, we aim to construct a physics-informed data-driven model, which happens to be linear, as most dynamic effects on the thermal mass of a building are indeed linear.

#### Modeling thermal zones

The evolution of the temperature  $T$  of a lumped mass in a building, for example a thermal zone, can be described by

$$mc_p \frac{dT}{dt} = \dot{Q}_{\text{amb}} + \dot{Q}_n + \dot{Q}_{\text{sol}} + \dot{Q}_{\text{occ}} + \dot{Q}_{\text{act}}, \quad (4.10)$$

where  $m$  and  $c_p$  denote the mass and specific heat capacity of the mass respectively. The  $\dot{Q}$  terms denote incoming and outgoing energy flows, i.e heat flows from and to ambient,

$\dot{Q}_{\text{amb}}$ , and neighboring zones,  $\dot{Q}_n$ , by solar irradiation,  $\dot{Q}_{\text{sol}}$ , by occupancy,  $\dot{Q}_{\text{occ}}$ , and by actuators  $\dot{Q}_{\text{act}}$  (such as radiators, air conditioning, etc.) respectively. The terms  $\dot{Q}_{\text{amb}}$  and  $\dot{Q}_n$  are linear functions of  $T$ , for example  $\dot{Q}_{\text{amb}} = \theta_{\text{amb}}(T_{\text{amb}} - T)$ , where  $\theta_{\text{amb}}$  is a constant;  $\dot{Q}_n$  is given by a sum of similar linear functions with  $T_{\text{amb}}$  replaced by the temperatures of the neighbouring zones. The influence of occupancy,  $\dot{Q}_{\text{occ}}$ , is part of ongoing research and is neglected for this particular model, partly because occupancy forecasts were not available in the case study. We will treat the remaining terms,  $\dot{Q}_{\text{sol}}$  and  $\dot{Q}_{\text{act}}$ , in the following.

### Modeling solar gains $\dot{Q}_{\text{sol}}$

A simple physics-based model for solar gains through windows is given by

$$\dot{Q}_{\text{sol}} = A_{\text{win}} \sin(\alpha - \alpha_0) \frac{\cos(\beta)}{\sin(\beta)} I_{\text{hor}}, \quad (4.11)$$

where  $A_{\text{win}}$  is the window surface area, and the angles  $\alpha$  and  $\beta$  denote the azimuth (i.e. the horizontal angle with respect to north) and elevation (i.e. the vertical angle with respect to earth's surface) of the sun respectively. The offset  $\alpha_0$  denotes the orientation of the window and  $I_{\text{hor}}$  is the horizontal global irradiation, which is an input commonly available from a weather forecast. Note that  $\sin(\alpha - \alpha_0)$  can become negative, which means that the direction of the irradiation through the surface is negative. For the given case of a window, we therefore take  $\max(0, \sin(\alpha - \alpha_0))$ . Equation (4.11) shows that the gains through solar irradiation are a non-linear function of the window orientation and the angles  $\alpha$  and  $\beta$ , which are themselves highly non-linear functions of time.

To embed  $\dot{Q}_{\text{sol}}$  in our model, we assume that the second term of (4.11),  $I_{\text{vert}(t)} = \frac{\cos(\beta)}{\sin(\beta)} I_{\text{hor}}$ , which denotes the irradiation on a vertical surface following the sun, is given as an input. This is a reasonable assumption as  $I_{\text{hor}}$  can be obtained from a weather forecast and  $\beta$  is a function of time and location. We model the first term,  $A_{\text{win}} \sin(\alpha - \alpha_0)$ , with  $\tau$  time varying coefficients  $[\theta_{\text{sol},1} \dots \theta_{\text{sol},\tau}]$ , where  $\tau$  represents one day. The solar gain can then be formulated as

$$\dot{Q}_{\text{sol}} = [\theta_{\text{sol},1} \dots \theta_{\text{sol},\tau}] [I_{\text{vert},1} \dots I_{\text{vert},\tau}]^T, \quad (4.12)$$

where  $[I_{\text{vert},1} \dots I_{\text{vert},\tau}]$  is a one-hot encoding (Brownlee, 2018) of  $I_{\text{vert}}$  with respect to discrete time-periods  $1, \dots, \tau$ :

$$I_{\text{vert},i} = \begin{cases} I_{\text{vert}(t)}, & \text{if } i \geq t > i + 1 \\ 0, & \text{otherwise.} \end{cases} \quad (4.13)$$

This way of modeling the solar irradiation creates  $\tau$  input variables that are zero during most times of the day, but are equal to  $I_{\text{vert}(t)}$  for fixed periods. For example, in the case of  $\tau=4$  with time periods of equivalent length,  $I_{\text{vert},1}$  attains  $I_{\text{vert}(t)}$  for the first six hours of the day, and is zero for all other times,  $I_{\text{vert},2}$  attains  $I_{\text{vert}(t)}$  for the period of 6 am to 12 pm, and zero otherwise, etc. The solar gains  $\dot{Q}_{\text{sol}}$  are now a linear function of the easy-to-obtain inputs  $I_{\text{vert},t1} \dots I_{\text{vert},\tau}$ . A validation of this modeling approach compared to the physical model of (4.11), on the data set later used in the case studies, is shown in Figure A.6 in the Appendix for  $\tau = 9$ . The coefficient of determination  $R^2$  is 0.96 for both the training and the testing set. Note that the number of one-hot encoded inputs  $\tau$  does not relate to the sampling time of the MPC; it only needs to be sufficiently high to reach a reasonable fit of  $\dot{Q}_{\text{sol}}$ .

### Modeling actuator gains $\dot{Q}_{\text{act}}$

For most relevant heating and cooling systems, such as radiators, floor heating and (neglecting the water and vapour content) air conditioning units, the energy transferred from an actuator to a thermal zone can be described by an equation of the form

$$\dot{Q}_{\text{act}} = \dot{m}_f c_{p,f} (T_{\text{sup}} - T_{\text{ret}}). \quad (4.14)$$

Here,  $\dot{m}_f$  and  $c_{p,f}$  denote respectively the mass flow and specific heat capacity of the fluid (usually water or air), and  $T_{\text{sup}}$  and  $T_{\text{ret}}$  denote supply and return temperatures.

There are a variety of options to model  $\dot{Q}_{\text{act}}$ . In an ideal case, it is directly accessible by measuring  $\dot{m}$ ,  $T_{\text{sup}}$  and  $T_{\text{ret}}$  and can be used directly in (4.10). Unfortunately, in many buildings measurements at this level are not available. In such cases, one has several options for inferring  $\dot{Q}$  from the available measurements. One option is by measuring the total energy consumption of a building and allocating portions of it based on design mass flows to individual rooms; this is the approach we follow for  $\tilde{\dot{Q}}_{\text{act},i}$  in Section 4.4. Another is to model  $\dot{Q}_{\text{act}}$  as a linear function of the mass flow or a valve position  $b$ :

$$\dot{Q}_{\text{act}} = \bar{\theta} \dot{m} \approx \theta b. \quad (4.15)$$

This assumes that  $(T_{\text{sup}} - T_{\text{ret}})$  is approximately constant, which is a reasonable assumption for example in the case of heating with a high supply temperature. In the case of using the valve position  $b$  and not directly the mass flow  $\dot{m}$ , the assumption that their relation is linear also has to be made.

A last option is suitable when both valve position  $b$  and  $T_{\text{sup}}$  are measured. The transferred energy then follows

$$\dot{Q}_{\text{act}} = \theta b (T_{\text{sup}} - T), \quad (4.16)$$

where  $T$  is the current temperature of the zone itself<sup>1</sup>. For air-based systems, the assumption  $T_{\text{ret}} = T$  is always reasonable. For water-based systems, the assumption holds if the heat transfer surfaces are large and the mass flows are low. Also, the assumption that  $\dot{m}$  and  $b$  have a linear relation has to hold again. All options assume that  $c_{p,f}$  is constant. In the case studies in Sections 4.5 and 4.6, we will explore several of these options for  $\dot{Q}_{\text{act}}$ .

### Positivity constraints for building dynamics

By substituting all  $\dot{Q}$  terms in (4.10) using (4.16) for  $\dot{Q}_{\text{act}}$ , and neglecting occupancy, the dynamics can be rewritten as

$$\begin{aligned}
 mc_p \frac{dT}{dt} = & \theta_{\text{amb}}(T_{\text{amb}} - T) + \theta_n(T_n - T) \\
 & + \left( \sum_{t_i=t_1}^{\tau} \theta_{\text{sol},t_i} I_{\text{vert},t_i} \right) + \theta_{\text{act}} b(T_{\text{sup}} - \bar{T}),
 \end{aligned} \tag{4.17}$$

where we assume a single neighbouring zone to simplify the notation; similar equations are obtained with the other options for modeling  $\dot{Q}_{\text{act}}$ . To avoid the bilinearity in the valve opening  $b$  and the zone temperature  $T$ , the zone temperature is replaced by an approximation  $\bar{T}$ , which can be obtained from the last measured room temperature for example. After performing Euler discretization, the discrete time thermal zone dynamics can be written as

$$\begin{aligned}
 T_{k+1} = & \left( 1 - \frac{\Delta t \theta_{\text{amb}}}{mc_p} - \frac{\Delta t \theta_n}{mc_p} \right) T_k + \left( \frac{\Delta t \theta_{\text{amb}}}{mc_p} \right) T_{\text{amb},k} + \left( \frac{\Delta t \theta_n}{mc_p} \right) T_{n,k} \\
 & + \left( \sum_{t_i=t_1}^{\tau} \left( \frac{\Delta t \theta_{\text{sol},t_i}}{mc_p} \right) I_{\text{vert},t_i} \right) + \left( \frac{\Delta t \theta_{\text{act}}}{mc_p} \right) b(T_{\text{sup}} - \bar{T})_k,
 \end{aligned} \tag{4.18}$$

where the subscripts  $k$  and  $k+1$  denote the current and subsequent time step respectively, and  $\Delta t$  is the sampling time. By reducing all constants in (4.18) to coefficients  $\tilde{\theta}$ , the expression can be further simplified to

$$\begin{aligned}
 T_{k+1} = & \tilde{\theta}_{\text{auto}} T_k + \tilde{\theta}_{\text{amb}} T_{\text{amb},k} + \tilde{\theta}_n T_{n,k} \\
 & + \left( \sum_{t_i=t_1}^{\tau} \tilde{\theta}_{\text{sol},t_i} I_{\text{vert},t_i} \right) + \tilde{\theta}_{\text{act}} b(T_{\text{sup}} - \bar{T})_k.
 \end{aligned} \tag{4.19}$$

---

<sup>1</sup>The input is therefore linear in the system state.

We emphasize that for a sufficiently small time step  $\Delta t$ , all coefficients  $\tilde{\theta}$  are positive since the masses, specific heat capacities and heat transfer coefficients are all positive.

The temperatures can be stacked in a state vector  $x$ , inputs  $(b(T_{\text{sup}} - \bar{T}))$  and disturbances  $(T_{\text{amb}}, I_{\text{vert}, t_i})$  can be stacked in the input and disturbance vectors  $u$  and  $d$ , and a conventional linear state space system of the form  $x_{k+1} = Ax_k + B_u u_k + B_d d_k$ ,  $y_{k+1} = Cx_k$  can be formulated, where the output vector  $y$  collects all the entries of  $x$  (zone temperatures) that can be measured. The influence of the hidden states can then be modeled implicitly by a convolution of the previous outputs  $y$  (i.e. the measured zone temperatures), inputs and disturbances. This results in an ARMAX model of the form

$$y_{k+1} = \Theta [y_k \dots y_{k-\delta} \ u_k \dots u_{k-\delta} \ d_k \dots d_{k-\delta}]^T, \quad (4.20)$$

which can be used for the dynamics in (4.1b). Here,  $\Theta$  is the vector of regression coefficients for each thermal zone and  $\delta$  determines the number of autoregressive and moving average steps.<sup>2</sup> Similar to the state space system, the elements of  $\Theta$  are positive and can be found through non-negative least squares regression.

### 4.3.4 Model properties

The presented models have different properties in terms of expressiveness and the complexity of the resulting MPC problem. The ARMAX model is limited to (4.1b) being linear in all variables. RF are generally more expressive, as they are piecewise constant in the disturbances and piecewise linear in the inputs. ICNN are also more expressive than ARMAX, being convex in the inputs, and even less restrictive for disturbances. Generally it could be expected that more expressive models lead to a more accurate representation of the system dynamics as long as they can be trained to optimality and there is enough data to avoid overfitting. This is studied in Section 4.6 below.

Compared to ICNN, ARMAX and RF are generally more lightweight in terms of online computation, when applied in the MPC problem (4.1). For example, if (4.1a) is quadratic and (4.1b) linear in the decision variables, the resulting problem is a QP, for which very efficient solvers are available. By contrast, ICNN lead to a general convex optimization problem, that require more generic methods like interior point or gradient decent. To get a convex problem, ICNN are limited to box constraints with only an upper bound when it comes to the output variables in (4.1c). The other two models allow general convex output constraints.



Figure 4.3: Rendering of the UMAR unit in NEST with both bedrooms marked.  
© Werner Sobek.




Figure 4.4: Heating system of the UMAR unit in NEST. The heating loop is connected to the central heating system via a heat exchanger. The heating and cooling panels in each room are controlled with individual on/off valves.

## 4.4 Test bed

We use the Urban Mining and Recycling (UMAR) unit (Kakkos et al., 2019; Heisel et al., 2019) of the NEST building (Richner et al., 2017) at Empa, Dübendorf, Switzerland as a test bed in our case studies. The unit, shown in Figures 2.2 and 4.3, is an apartment built to demonstrate the circular economy in the building construction industry and is constructed from recycled material or material that can be recycled completely after dismantling the unit. It comprises two bedrooms and one living/kitchen area with large south-east facing windows, two bathrooms, an entrance area and a technical room. The two bedrooms have identical floor plans and furniture. The apartment is considered to be a “living lab” and is occupied by two persons, who live there.

The unit is equipped with a water-based heating and cooling system (Figure 4.4). The

<sup>2</sup>Generally,  $\delta$  can be different for  $y$  and  $u, d$ .

system is connected to the NEST heating grid via a heat exchanger and the total energy consumption of the unit,  $\dot{Q}_{\text{UMAR}}$ , is measured with the help of an Energy Valve™ (Belimo, 2021). Each room is equipped with at least one heating/cooling ceiling panel, which is controlled by individual on/off valves  $b_i$ . There is a central pump that provides pressure, as long as one of the valves is open. The supply temperature  $T_{\text{sup}}$  is regulated by adjusting the mass flow on the NEST heating grid side of the heat exchanger. The same system is used for cooling during warm days through a second heat exchanger connected to the NEST cooling grid. As the pump delivers a constant pressure and the valves are either fully open or closed, the design mass flows<sup>3</sup> for each heating/cooling panel  $\dot{m}_i$  can be used to calculate the amount of energy transferred to each room  $\tilde{Q}_{\text{act},i}$  through

$$\tilde{Q}_{\text{act},i} = \dot{Q}_{\text{UMAR}} \frac{\dot{m}_i b_i}{\sum_{j \in \mathcal{R}} \dot{m}_j b_j}, \quad (4.21)$$

where  $\mathcal{R}$  denotes the set of all rooms in the unit. All measurements, including the individual room temperatures are stored in an SQL database with a sampling time of one minute. A weather forecast, sampled in one-hour periods, for the ambient temperature and global solar irradiation is available from the national weather service, Meteo Swiss. All actuators can be controlled via OPC-UA software clients (Leitner and Mahnke, 2006). During standard operation, the room temperature is controlled by a hysteresis baseline controller that regulates the room temperature between an occupant-decided set point and 1 °C below in the cooling case, and above in the heating case. An example of a temperature trajectory of bedroom 2 and the corresponding control input with the baseline controller is shown in Figure A.7 in the Appendix.

## 4.5 Experiment case study

Over the course of two years, we have implemented MPC, based on RF, ICNN and physics-informed ARMAX models, for 156 full days of experiments in the two bedrooms of the UMAR unit. In the following, we will discuss the model configuration and training, present example closed-loop trajectories of the controllers, and analyse the energy consumption compared to the base-line controller. For each controller type a different model was trained for each bedroom.

---

<sup>3</sup>We introduce this way of calculating the transferred energy in Section 4.4 rather than 4.3, as the design mass flows of each individual zone are commonly not known in a building and specific to this case study. They are specified by the heating/cooling system installer.



### 4.5.1 Model and controller configuration

All models use the same training data set of 2018-05-23 to 2019-05-28 (370 days), generated during normal operation of the building with the baseline hysteresis controller. In this subsection, we define which inputs and outputs each model uses, how the model hyperparameters were determined, and how many degrees of freedom each model has.

To configure the RF model, extensive feature engineering and hyperparameter tuning was conducted (Bünning et al., 2020b). For the feature engineering, domain knowledge was used to pre-select certain model inputs and disregard others. With these features, models were then trained on the first 70% of the data and tested on the remaining data. As a result of feature engineering, it was found that predicting room temperature differences  $\Delta x$  (or  $\Delta T$ ) leads to better model accuracy than directly predicting room temperatures. This has no effect on the convexity of the MPC problem (4.1). As model inputs for the forests, autoregressive terms of  $\Delta x$ , of temperature differences between rooms and neighboring rooms, of window opening times, of the horizontal solar irradiation, and of the ambient temperature were chosen. We also use the time of the day and month, encoded as cosine and sine functions<sup>4</sup>, to capture any time-related influences on the room temperatures (for example, occupants entering the rooms every day at the same time) and to identify relations between time and global solar irradiation (similar to (4.12)). For the linear regression in the forest leaves, moving average terms of the control input  $\tilde{Q}_{\text{act},i}$  are used. The number of autoregressive and moving-average steps was set to 16.<sup>5</sup> The sampling time was set to 30 minutes based on preliminary numerical studies.

The hyperparameters of a random forest are mainly the number of trees per forest and the minimum number of samples per leaf. To find suitable hyperparameters, we applied line searching, where each parameter is varied while keeping the others constant. The procedure is not iterative, i.e. each parameter is just updated once. The resulting model has 200 trees per forest and a minimum of 200 samples per leaf.

The degrees of freedom of a random forest with linear regression grows with the number of training samples  $S$ , due to the automatic scaling through the minimum number of samples per leaf, and approximately follows  $3\frac{S}{200}$  (with 200 being the number of samples per leaf). For a sampling time of 30 minutes, this amounts to 270 for the full training set. The models are implemented with *Scikit-learn* (Pedregosa et al., 2011) in *Python 3*.

The configuration of the ICNN models is described in (Bünning et al., 2021b) and in more detail in (Schalbetter, 2020). For feature engineering, a k-fold cross validation with

<sup>4</sup>Both cosine and sine are used, to have distinct inputs for each time, as sine and cosine functions have the same function value twice per period (Bescond, 2020). We can also fit any phase offset with a linear parametrization.

<sup>5</sup>This number might seem excessive, however, additional inputs do not affect the fitting process of regression trees negatively, as unimportant inputs do not get chosen as splitting/partitioning variables.

$k=12$  was performed on the entire data set, with 9 folds being used for training and 3 for validation. After this, the networks were chosen to predict the change of room temperatures instead of directly predicting temperatures. For PICNN, as non-convex inputs, autoregressive terms of the solar irradiation (with three delay steps), the temperature difference with respect to neighboring rooms (one delay step) and the sine-encoded time of day (but not time of year) were chosen (one delay step). As convex inputs, autoregressive terms of the change of room temperature (with three delay steps), the temperature difference between room and ambient (one delay step), and the heating/cooling control input were chosen (one delay step). In the case of FICNN, all the above features are convex inputs. The sampling time was also decided on with cross validation and is 20 minutes for predictions up to one hour and 180 minutes beyond (Schalbetter, 2020). To allow the network output  $z_L$  (i.e. the change of room temperature) to be negative, we use a shifted ReLU function,  $g_i(x) = \max(x, 0) - \beta$ , as the final activation function in the output layer of the networks instead of a regular one. Here,  $\beta$  is a hyperparameter of the network.

The other hyperparameters of an ICNN are the training method, the step-size of the training method, number of training epochs, nodes per layer and number of layers. They were optimized with line searches applied to the  $k$ -fold cross validation. The resulting networks have approximately 1000 degrees of freedom (i.e. parameters to fit). The models are implemented with *Keras* (Chollet, 2018) in *Python 3*.

The inputs for the ARMAX model directly follow from the physical structure described in Section 4.3. They comprise autoregressive terms for the room temperature<sup>6</sup>, moving average terms for neighboring zones, the ambient temperature, the one-hot encoded global solar irradiation and for the control input. The only hyperparameters to be tuned are  $\tau$ , the number of one-hot inputs of the solar irradiation, and  $\delta$ , the number of autoregressive and moving average terms. These were chosen to be  $\tau = 9$  and  $\delta = 3$  after preliminary experiments. Our ARMAX models have  $(\delta + 1)(4 + \tau)$  degrees of freedom, i.e. 52 for our configuration. The models are implemented with *Scikit-learn* (Pedregosa et al., 2011) in *Python 3*.

All models are applied as the model for the building dynamics (4.1b) in the MPC optimization problem (4.1). After specifying the cost function and constraints, this leads to the optimization problem

$$\min_{u, y, \epsilon} \sum_{k=0}^{N-1} (u_k R u_k + Q u_k + \lambda \epsilon_{k+1}) \quad (4.22a)$$

$$\text{s.t.} \quad y_{k+1} = f(y_{k-}, u_{k-}, d_{k-}) \quad (4.22b)$$

$$y_{\min} - \epsilon_{k+1} \leq y_{k+1} \leq y_{\max} + \epsilon_{k+1} \quad (4.22c)$$

---

<sup>6</sup>For linear models, the choices of using absolute temperatures or temperature differences both lead to the same model accuracy.

$$\epsilon_{k+1} \geq 0 \tag{4.22d}$$

$$u_{\min} \leq u_k \leq u_{\max} \tag{4.22e}$$

$$\forall k \in [0, \dots, N - 1],$$

where  $\epsilon$  is a slack variable introduced to ensure feasibility. A quadratic weight for the control input  $R$ , a linear weight for the control input  $Q$ , and a linear weight  $\lambda$  for the comfort slack variable are used in the cost function; the weights are specified below for each experiment. The comfort constraints  $y_{\min}$  and  $y_{\max}$  are time varying and will also be reported with the results. The control input is applied to the valves of the UMAR unit with pulse-width modulation (PWM). The limits for the control input are therefore  $u_{\min} = 0$  and  $u_{\max} = 1$ . The resulting problem is a Quadratic Program in the case of the ARMAX and RF models (see discussion in Sections 4.3.1 and 4.3.4), which we solve with the QP solver of *CVXOPT* (Andersen et al., 2004) in *Python 3*. For the ICNN, under the absence of the lower output constraint, the problem results in a convex problem without direct access to the function derivatives. We solve it with the *COBYLA* (Powell, 1994) solver of *SciPy* (Virtanen et al., 2020) in *Python 3*.

## 4.5.2 Example closed-loop experiments

To compare the closed-loop behaviour of the ARMAX controller and an FICNN controller, we conducted a cooling experiment on NEST with the ARMAX model (controller properties:  $N=7\text{h}$ ,  $R=1$ ,  $Q=0$ ,  $\lambda=100$ ,  $\delta=7$ ,  $\Delta t=30$  min, actuation: valve opening) applied to bedroom 1 and a FICNN (controller properties:  $N=7\text{h}$ ,  $R=1$ ,  $Q=0$ ,  $\lambda=100$ ) applied to bedroom 2. The results in Figure 4.5 show that both controllers keep the temperature within the comfort constraints most of the time and exploit the relaxed constraints during the day to save energy. However, the controller using the ARMAX model is less conservative. It keeps the room temperature closer to the upper comfort constraint during the night, and meets the lowered comfort constraint at 22:00 just in time or violates it by a fraction of a degree for a short period of time. Day 2020-08-19 is characteristic of the different behaviour of the two controllers. Here, the FICNN-controller applies control action during the second half of the day although the temperature is already relatively low, (which even leads to slight violations of the lower comfort constraints a few hours later,) while the ARMAX-controller slightly violates the upper comfort constraint. Experiments with the PICNN showed similar results to those conducted with the FICNN (Data omitted in the interest of space).

To compare the behaviour of the RF controller in a similar setting, we performed a cooling experiment with an RF controller applied to bedroom 1 (controller properties:  $N=6\text{h}$ ,  $R=1$ ,  $Q=0$ ,  $\lambda=100$ ); to investigate the effect of the cost weights we also applied the same controller to bedroom 2 with controller properties  $N=6\text{h}$ ,  $R=100$ ,  $Q=0$ ,  $\lambda=100$ . Figure 4.6 demonstrates that the behaviour is very similar to the one observed with the




Figure 4.5: Cooling experiment with ARMAX in bedroom 1 (blue) and FICNN in bedroom 2 (orange). (a): Temperature in the two bedrooms. The comfort bounds are shown in dashed black. In ①, ③ and ④ the connection between the controller and actuators was lost for a short time. In ②, the otherwise closed window blinds were automatically opened due to strong wind, which lead to the system not being able to reject the solar gains, even at maximum cooling power. (b): Relative control input, i.e. the fraction of time where the maximum control input is applied during one control step. (c): Measured ambient temperature at the experiment site. (d): Global solar irradiation at the experiment site.

ARMAX model. The room temperature is kept close to the upper comfort constraint during the night, the controller exploits the relaxed comfort constraints during the day to save energy, and starts cooling early enough to meet the lowered comfort constraints at 22:00. The difference in relative weighting between costs for control input and constraint violations does not seem to significantly influence behaviour.

In general, the experiments demonstrate that all controllers have reasonable behaviour, with the ICNN-based controllers being more conservative compared to the ARMAX and RF-based controller. This is possibly due to underestimating the influence of the control input or overestimating the thermal capacity of the system. Although this issue is more pronounced for the ICNN, it is also visible for RF and ARMAX to a lesser extent during times where the comfort constraints are relaxed but the controller is aiming for the lowered upper comfort constraint at 22.00, for example at ① in Figure 4.6. Here, the controllers often apply a high control input in one step, which cools down the room temperature more than necessary, and then do not apply any control input in the consecutive time step, which lets the room temperature rise again. This is likely a result of the training data being correlated due to the underlying feedback controller: the




Figure 4.6: Cooling experiment with RF in bedroom 1 (blue) and bedroom 2 (orange) with different weights on the control input cost. (a): Temperature in the two bedrooms. The comfort bounds are shown in dashed black. (b): Relative control input, i.e. the fraction of time where the maximum control input is applied during one control step. (c): Measured ambient temperature at the experiment site. (d): Global solar irradiation at the experiment site.

effects of control input and disturbance on the room temperature cancel out. Besides the obvious solution of generating training data with an uncorrelated input signal, tracking the predicted temperature trajectory with a lower-level feedback controller instead of directly applying the optimized control input could mitigate the issue.

Similar experiments in UMAR have been conducted for the heating case with the RF and the ARMAX controller. (We have not conducted heating experiments with the ICNN models due to the MPC problem not being convex at the lower comfort constraint (Bünning et al., 2021b).) The general behaviour of the controllers in the heating experiments is similar to the one observed in the cooling experiments. The results therefore do not add anything new to the discussion, other than the observation that predictive control in general also works in the heating case with both model types. We therefore refer the interested reader to the linked data repository at the end of the chapter for data on these experiments.

### 4.5.3 Energy consumption

We compare the energy consumption of the MPC-based controllers to the baseline controller on the basis of the concept of Heating Degree Solar Days (HSDS) and Cooling

Degree Solar Days (CDS), explained in the following.

### HDD and CDD

Heating Degree Days (without *Solar*) are commonly used to quantify the energy consumption of a building as a function of the ambient conditions (ScienceDirect, 2021). Here, a base temperature  $T_{\text{amb},b}$  is defined (usually assumed to be the ambient temperature at which no heating is necessary), and it is assumed that the daily heating energy consumption of a building is proportional to the difference between the daily average ambient temperature  $\bar{T}_{\text{amb}}$  and the base temperature:

$$Q_{\text{hea}} = \theta_{\text{HDD}}(T_{\text{amb},b} - \bar{T}_{\text{amb}}) = \theta_{\text{HDD}}\text{HDD}. \quad (4.23)$$

The difference between the base temperature and the daily average temperature is called Heating Degree Days (HDD). The coefficient  $\theta_{\text{HDD}}$  can be found with linear regression. Similarly a model with Cooling Degree Days (CDD) can be defined to quantify the cooling consumption of a building based on a separate base temperature  $T_{\text{amb},b}$ . (Depending on the selection of the different base temperatures, the same day can have Heating Degree Days and Cooling Degree Days). While this expression might be a good assumption for most common buildings, it might not be for buildings with large windows, such as the UMAR apartment, where the solar irradiation is a main driver of the dynamics. We therefore add the global solar irradiation  $I_{\text{hor}}$  with a new regression coefficient  $\theta_{\text{sol}}$  to the equation,

$$Q_{\text{hea}} = \theta_{\text{HDD}}\text{HDD} + \theta_{\text{sol}}I_{\text{hor}} + c, \quad (4.24)$$

and add a constant  $c$  to omit the dependence on the base temperature  $T_{\text{amb},b}$ .<sup>7</sup> Finally, we divide the regression coefficients by each other and define the Heating Degree Solar Day (HSD):

$$Q_{\text{hea}} = \theta_{\text{HDD}}\left(\text{HDD} + \frac{\theta_{\text{sol}}}{\theta_{\text{HDD}}}I_{\text{hor}}\right) + c = \theta_{\text{HDD}}\text{HSD} + c. \quad (4.25)$$

The concept of Cooling Degree Solar Days (CSD) works accordingly, approximating the cooling demand  $Q_{\text{coo}}$  of a building.

---

<sup>7</sup>If the regression chooses to set  $c = -\theta_{\text{HDD}}T_{\text{amb},b}$ , the function effectively is not dependent on  $T_{\text{amb},b}$  any more. We could also write (4.24) directly in terms of  $\bar{T}_{\text{amb}}$ , but choose not to, to maintain the relation to the HDD concept.

## Experiment results

We have evaluated measurement data spanning 2.5 years (2018-06-01 to 2021-02-14) for the analysis of the energy consumption with the MPC controllers compared to the baseline controller (example in Figure A.7 in the Appendix).


Figure 4.7 (a) shows the heating energy consumption as a function of the HDS D for bedroom 1. Each blue dot represents a day under standard operation with the baseline controller and each orange dot represents a day with an MPC controller. Days where the average room temperature is outside the range of 21 °C to 27 °C are excluded from the analysis for a fair comparison, as these conditions are often results of the heating/cooling system not working correctly, leading to unrealistically low/high energy consumption. The lines represent the HDS D regressions. Note, that we have lumped all different MPC models (ARMAX, RF and ICNN)<sup>8</sup> here in orange, as all controllers have shown reasonable control performance in the individual experiments. As the experimental data sets for the individual methods are small, we refrain from distinguishing quantitatively between the modeling methods. However, as we use a different marker for each method, it can be seen that qualitatively, the methods perform similarly.

There is a clear trend visible confirming that the MPC controllers consume less heating energy than the baseline controller. At 5 HDS D the reduction is approximately 41% while at 15 HDS D it is 29%. Besides anticipating “free of cost” heat gains from solar irradiation and other environmental factors, the MPC controllers of course save a significant amount of energy by exploiting relaxed comfort constraints. However, we note that the energy savings do not stem from a generally substantially lower room temperature set point, or the violation of constraints (as could be seen in Section 4.5.2). As Figure A.8 in the Appendix shows<sup>9</sup>, the median daily room temperature difference between MPC and standard operation is 0.18 °C. The same trend of MPC consuming significantly less energy compared to the baseline controller is visible in the cooling experiments in bedroom 1, shown in Figure 4.7 (b). MPC consumes 33% less cooling energy at 5 CDS D and 28% less at 15 CDS D.


The trend is slightly less pronounced in the heating experiments in bedroom 2, which are shown in Figure 4.7 (c). Here, the difference between the MPC controllers and the baseline controller is smaller. This behaviour can be explained by the better insulation of bedroom 2. While bedroom 1 has a window surface and a wall surface connected to the ambient, the wall surface of bedroom 2 is adjacent to another unit. This results in less temperature loss during the night, which means that the MPC controller can exploit the relaxed comfort constraint less. The result indicates that the potential for energy

<sup>8</sup>The *ARX+window* label denotes a model where eq. (4.11) is calculated based on the window dimensions and orientations and directly used as a model input instead of eq. (4.12).


<sup>9</sup>Data on the cooling case in bedroom 1 and both cases in bedroom 2 is available in the linked data repository, see Section *Data Availability*. These data show a similar trend.




(a) Heating case bedroom 1.



(b) Cooling case bedroom 1.



(c) Heating case bedroom 2.



(d) Cooling case bedroom 2.

Figure 4.7: Heating energy of MPC methods (orange) and baseline (blue) controller as a function of HDS and CDS. Each sample represents one day of experiment, the solid lines show the HDS and CDS regressions of these samples. (a): Heating case bedroom 1 with 41 samples for MPC, 258 samples for baseline controller. (b): Cooling case bedroom 1 with 39 samples for MPC, 230 samples for baseline controller. (c): Heating case bedroom 2 with 41 samples for MPC, 184 samples for baseline controller. (d): Cooling case bedroom 2 with 35 samples for MPC, 206 samples for baseline controller.



savings through MPC depends on the level of insulation of the controlled building. For the cooling experiments in bedroom 2, which are shown in Figure 4.7 (d), MPC again saves a significant amount of energy compared to the baseline controller. Here, the solar irradiation is the significant disturbance working against the control input, and the window surface area is the same for both bedrooms, which means that relaxed comfort constraints can be exploited in both bedrooms.

When the HDSD and CDSO input data of the entire 2.5 years of measurement data is applied as an input to the linear regression models<sup>10</sup> in Figure 4.7, the MPC controllers on average save 33% and 26% of heating energy in bedrooms 1 and 2 respectively, and 32% and 49% of cooling energy in bedrooms 1 and 2 respectively compared to the baseline controller. Note that, given the deviations of the individual experiments from the regressions in Figure 4.7, this is only a rough estimate. However, it demonstrates that MPC generally saves a significant amount of heating and cooling energy in our experiments, without a significant difference in room temperatures or constraint violations.

#### 4.5.4 Computational performance

Table 4.1 summarises the computational requirements for the MPC controllers tested with the different models of this chapter on a single optimization with fixed initial conditions and disturbance forecasts. The optimization problems were solved on an Intel(R) Core(TM) i7-7500U CPU with 2.7 GHz, and 16 GB of memory. The analysis is done for  $N=6h$ ,  $R=1$ ,  $Q=0$ ,  $\lambda=100$ .

It can be seen that the ARMAX controller outperforms the other two in terms of solving time and memory usage. While the RF problem is also a QP, the solving time is longer because the parameters for the optimization problem need to be found from the forest on the basis of the non-controllable inputs  $X_d$  first. The PICNN has a longer solving time as the optimization problem is more difficult to solve. This is because, even though the problem is convex (without the presence of lower comfort constraints), we do not have access to analytical gradients and the solver has to numerically approximate them. FICNN show similar results. The memory requirements grow with the complexity of the models. While all solving times are acceptable for the considered residential building case, it is to be expected that RF and ARMAX controllers scale better due to their linear dynamics, which could be beneficial for larger scale buildings, or the reserve application discussed in Chapter 5. We note that all models and solvers are based on open-source libraries, which are not optimized for our specific purpose.

---

<sup>10</sup>This corresponds to a situation where either the baseline controller or the MPC would have been run for the entire duration of 2.5 years.

Table 4.1: Computational resources for MPC controllers with varying models.

Model	Solving time	Memory usage	Software (Python 3)
ARMAX	0.2 s	36 MB	Scikit-learn, CVXOPT
RF	1.5 s	68 MB	Scikit-learn, CVXOPT
PICNN	3.0 s	158 MB	Keras, SciPy

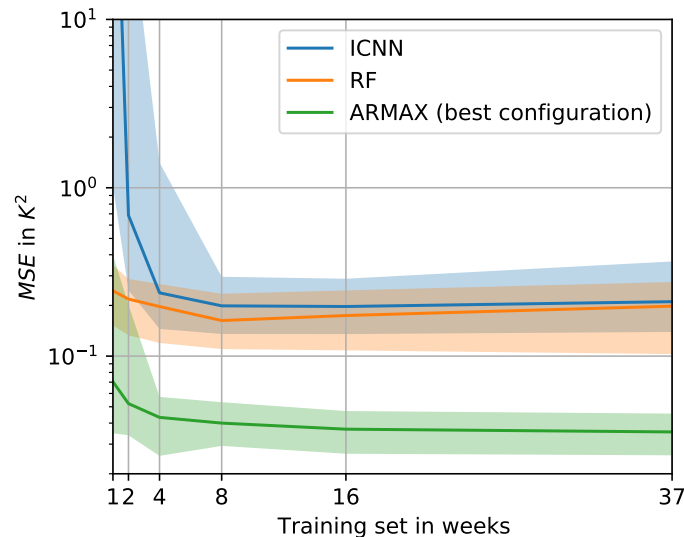


Figure 4.8: Mean Squared Error for 1-hour open loop prediction with ICNN (blue), RF (orange) and ARMAX (green) models. The solid lines depict the median achieved MSE, and the shaded areas depict the 16% and 84% percentiles.

## 4.6 Numerical case study

Our experiments have shown that MPC with all the considered models is generally suitable for building control in practice. However, to minimize costs and time of commissioning, the amount of historical measurement data required to train the model is a considerable factor for choosing a model for MPC. To address this point, we compare here the sample efficiency of the ARMAX, RF and ICNN models.

For this, the training data from 2018-05-23 to 2019-05-28 for bedroom 1 is divided into weekly folds and a k-fold cross validation is performed, where a randomly selected subset of the data is used for training and the rest is used for model validation in terms of the mean squared error (MSE) for a one-hour open loop prediction. The numerical experiment is repeated 100 times to account for different training and validation data. Figure 4.8 shows the result for a PICNN, RF (both tuned for this particular bedroom), and an ARMAX model with three lag terms<sup>11</sup> using a positivity constraint on the ele-

<sup>11</sup>Our validation experiments suggest that including more than three lag terms does not improve the accuracy of the predictions of the ARMAX model in this case.




Figure 4.9: Mean Squared Error for 1-hour open loop prediction with ARMAX model and varying control inputs. The solid lines depict the median achieved MSE, and the corresponding coloured shadows depict the 16% and 84% percentiles.

ments of  $\Theta$  in (4.20). We do not show results of FICNN as they are very similar to those of PICNN. All models use  $\tilde{Q}_{act,i}$  (Equation (4.21)) as the control input.

It can be seen that the ICNN performs poorly for a single week and two weeks of training data, and significantly improves after four weeks. This is most likely due to the large number of parameters that need to be fitted. In contrast, the RF already performs relatively well for a single week of training data but does not improve much when more training data is available. The good performance for a small training data set is due to the automatic scaling of RF through the minimum number of samples per leaf. The median of the ARMAX model outperforms both the ICNN and RF models for all sizes of training data. This could be expected for small training data sets as the physical priors lead to strong regularization. However, it can also be seen that the ICNN and RF models converge to the same median MSE and similar variance, while the ARMAX model's median converges to a much lower MSE and less variance. This indicates that using physical model inputs gives the ARMAX model domain knowledge that the ML models cannot find by themselves (for example the solar irradiation through windows as a function of time and global solar irradiation), even if abundant training data is available.

The higher variance of the prediction error of the ARMAX model for low-sized training sets, compared to the RF model, is most likely an effect of the one-hot encoding of  $I_{vert}(t)$  (the vertical solar irradiation): for small training sets it can happen that the set does not include samples from the summer, which means that  $I_{vert,1}$  and  $I_{vert,\tau}$  are always

zero in the training set because the sun rises late and sets early<sup>12</sup>. Accurate coefficients can therefore not be found during training. If samples from the summer are included in the validation set, this leads to large prediction errors. For real applications, the issue is likely less pronounced, if the models are updated regularly. The issue could also be addressed through further regularization, e.g. with the Lasso method (Tibshirani, 1996).

For practical implementations of data-driven MPC, it is important to know which quantities need to be measured in a building. We therefore compare ARMAX models with various choices for the control input, as described in Section 4.3.3, in Figure 4.9. The first model (blue) uses the considered *true* energy input  $\tilde{Q}_{\text{act}}$  (Equation (4.21)), the second one (orange) only uses the valve opening  $b$  (Equation (4.15)), and the third one (green) uses the approximation via valve opening, supply temperature  $T_{\text{sup}}$  and a previously measured room temperature  $\bar{T}$  (Equation (4.16)). The last model (red) additionally constrains all regression coefficients  $\Theta$  in (4.20) to be non-negative. The first three models perform similarly well, which indicates that using just the valve opening to model the control input could be sufficient in practical cases - an observation that is also supported by our experiments. Measuring mass flows and supply temperatures gives no visible performance advantage, which simplifies practical implementation. The red result in Figure 4.9 demonstrates that the positivity constraint on  $\Theta$  significantly benefits the sample efficiency. This is especially important for practical implementation, as it reduces controller commissioning time. The model variance and the median MSE are also reduced.

## 4.7 Conclusion

In this chapter, we have compared physics-informed ARMAX models to Machine Learning models based on RF and ICNN in the domain of MPC for building climate control in experiments and in numerical case studies.

It was shown that MPC with all three models, RF, ICNN and ARMAX, generally delivers good control performance. Moreover, in all heating and cooling cases, MPC achieves significant energy savings compared to the baseline controllers. Our results also suggest that the physics-informed ARMAX models outperform the RF and ICNN models in terms of online computational requirements and offline training sample efficiency, which means that good models can be extracted from less data. The latter finding suggests that the physics-based inputs and constraints give the model an information prior, which cannot be found by the ML methods themselves, even if abundant training data is available. The increased expressiveness of ML-based models therefore does not seem to add any benefits in this case.

---

<sup>12</sup>In our implementation  $I_{\text{vert},0}$  covers the night-time, therefore  $I_{\text{vert},\tau}$  (late evening) and  $I_{\text{vert},1}$  (early morning) may or may not be non-zero depending on the season.

Considering the larger scope of this thesis, this result is interesting when compared to the prediction results of the demand forecasting models of Chapter 2, where even the uncorrected Machine Learning methods outperformed the resistor-capacitor models (which effectively also impose a physical prior on the model). The difference in performance could be explained with the fact that the demand forecasting methods have to predict the closed-loop behaviour of the considered system, which is, although dominated by the dynamics of the considered buildings and their ambient conditions, also dependent on the underlying specific controller behaviour. Depending on the controller type, this behaviour can also be time dependent for example and ML methods appear to be more suitable to capture it.

For modeling building dynamics, the next logical step is to apply the physics-informed inputs to the ML methods and to find ways to enforce physics-based constraints with ICNN and RF. Physics-informed ML methods could potentially be interesting as soon as non-linearities are added directly to the control problem, instead of being treated on a lower control level. Moreover, experiments with the physics-informed ARMAX models should be conducted on large-scale buildings. This is currently under investigation with an industry partner.

## Data Availability

The data presented in this chapter, and data from additional experiments is available under the DOI <https://doi.org/10.3929/ethz-b-000496285>.



# Combining electric reserves with data-driven building control

As discussed in Chapter 3, district heating and cooling systems, supplied by large-scale heat pumps and chillers, and their connected buildings are candidates for providing electrical reserves to the electricity grid. In Chapter 3, we have presented a method to offer electrical reserves with such systems without the need of exploiting the thermal inertia of the connected buildings. This is necessary if buildings do not want to or cannot participate in the scheme, for example due to privacy concerns or hardware limitations. Other buildings, however, might be keen and able to participate, and could potentially increase the reserves offered by allowing their thermal inertia to be exploited. The methods for data-driven modeling of building thermal dynamics presented in Chapter 4 can be used to integrate them into the reserve provision scheme.

In this chapter, we therefore present a central control scheme based on Robust Model Predictive Control (MPC) with affine policies to offer frequency reserves with a district system, where we exploit the thermal inertia of buffer storage tanks and a subset of the connected buildings. We leverage the data-driven model generation methods of Chapter 4 to overcome the problem of physics-based building modeling being tedious because of each building being individual. In a numerical case study based on one-year historical data of a real system, we compare the approach to the configuration of Chapter 3, where only the buffer storage is used for flexibility. We demonstrate that the reserves offered are increased substantially if the inertia of a subset of the connected buildings is exploited too. Furthermore, we validate the control approach in a first-of-its-kind experiment on the actual system, where we show that while reserves are offered by the district system, the comfort in the connected buildings is maintained at all times.

## 5.1 Introduction

Buildings with electrified heating and cooling systems are potential candidates for providing electrical reserves on the demand side because their thermal inertia allows loads

to be shifted in time (Junker et al., 2018; Kathirgamanathan et al., 2020; Tang and Wang, 2021). However, due to minimum capacity requirements, individual buildings are often not allowed to participate in reserve programs. Instead, their loads have to be aggregated (Alahäivälä et al., 2017; Oldewurtel et al., 2013; Geidl et al., 2017). Besides aggregating individual buildings, entire district heating and cooling systems could potentially be used to provide reserves, if heating or cooling is supplied with large-scale heat pumps (Lygnerud et al., 2021; Harild Rasmussen et al., 2021), as discussed in Chapter 3. This possibility is facilitated by the trend towards low-temperature heating and high-temperature cooling networks (Wirtz et al., 2020; Lund et al., 2021) as these allow heat pumps to be operated at higher coefficients of performance, making them a cost-effective alternative to boilers or combined heat and power units. While some authors discuss the general potential of offering electrical reserves with electric district heating systems (Lund et al., 2018), others develop methods to estimate this potential for concrete cases (Xu et al., 2020). There are also efforts to develop dynamical models for individual components of such systems, to simulate their operation (Pagh Nielsen and Sørensen, 2021). However, according to reviews (Hennessy et al., 2019; Vandermeulen et al., 2018), there is a lack of studies investigating the control of such configurations, both in simulation and in application.

In Chapter 3, we have presented and validated a control approach to offer reserves on the day-ahead frequency regulation market with a heat pump serving a small district heating system. The Robust MPC exploits the thermal inertia of a buffer storage tank and increases the reserves offered by using affine policies on the uncertainty introduced by the frequency regulation signal and other disturbances. The connected apartments were only treated as a disturbance to the buffer and are not included in the dynamics of the optimization problem. This configuration addresses the situation where connected buildings cannot or do not want to participate in the reserve scheme, for example due to missing infrastructure or privacy concerns.

In addition, this choice was made because compared to easy to obtain building heating demand forecasts (Zhao and Magoulès, 2012; Wang and Srinivasan, 2017; Fouquier et al., 2013; Suganthi and Samuel, 2012; Cholewa et al., 2021), modeling building dynamics from first principles is considered to be tedious and cost-prohibitive for industrial application because buildings are individual (Sturzenegger et al., 2016). However, in Chapter 4 we have shown that even simple models based on physics-constrained linear regression are suitable for predictive control of building thermal dynamics.

### **5.1.1 Contribution**

In this chapter, we therefore present an extended control scheme that leverages these models to offer frequency reserves with a district heating and cooling network supplied by a central heat pump, which exploits both the thermal inertia of buffer storage tanks



and of a subset of the buildings connected to the district system. By controlling a subset of the buildings, we seek to increase the reserves compared to the case where only the thermal flexibility of a buffer storage is used. This configuration has relevance for practical applications as it resembles the situation where a subset of customers connected to the district heating/cooling system is willing to give control authority to the district heating/cooling provider, while others are not. The presented control scheme is a two-level Robust MPC scheme<sup>1</sup> that uses affine policies to reduce the effect of uncertain disturbances. The controlled buildings are modeled with a data-driven approach based on positive linear regression (as discussed in Chapter 4), the heating demand of the other connected buildings is forecast (as discussed in Chapter 2), while the heat pump and storage tanks are modeled based on physics. The approach is applied in a numerical case study based on measurement data spanning one year of the NEST demonstrator (Richner et al., 2017). We show that by exploiting the thermal inertia of a subset of connected buildings, the operational costs can be lowered and the reserves offered can be increased by a factor of 2.7 compared to the case where only the thermal inertia of the buffer storage is used. In a first-of-its kind experiment on the actual district heating and cooling system with connected apartments and office units, it is shown that the control scheme is viable for practical applications, i.e. reserves are offered while the comfort is maintained in all connected buildings.

### 5.1.2 Structure

The remainder of the chapter is structured as follows. In Section 5.2, we present the reserve provision scheme, the system under consideration, and the architecture of the control scheme. In Section 5.3, we first discuss the models that describe the energy system and then present Levels 1 and 2 of the Robust MPC control scheme. In Section 5.4, we present the considered case study and demonstrate the potential of the control approach in a numerical case study on the basis of one year of historical measurement data. In Section 5.5, we show the validation of the control scheme in an experiment on the demonstrator and discuss the results. In Section 5.6, we discuss some limitations of the study and we conclude the chapter in Section 5.7.

---

<sup>1</sup>Compared to Chapter 3, we only treat the first two control layers here, as the lower level controllers for heat pumps and buildings have been discussed in Chapters 3 and 4.

## 5.2 Problem statement

### 5.2.1 Reserve provision scheme

The considered day-ahead frequency regulation scheme has been presented in Chapter 3, but we revisit it here for convenience. It is inspired by the RegD product of PJM (*Pennsylvania, New Jersey, and Maryland*), a transmission system operator in the U.S.. The reserve provider, i.e. the district heating/cooling operator in this case, sends an offer of non-negative reserves to the TSO at the beginning of the day. The offer is divided in intervals of 15-minutes for the next 24 hours, thus  $r \in \mathbb{R}_+^{96}$ . When the offered reserves are due on the following day, the reserve provider is allowed to adapt their electric base load  $u_{\text{el},k}^0$  every 15 minutes, i.e. every time step  $k$ . The heat pump should then track the electric load

$$u_{\text{el},k}(\tau) = u_{\text{el},k}^0 + w(\tau)r_k, \quad (5.1)$$

where  $r_k$  is the  $k^{\text{th}}$  element of the vector of reserves offered  $r$ , and the regulation signal  $w(\tau) \in [-1, 1]$  is updated by the TSO every 2 seconds. Note that the offered reserves are effectively symmetric because  $w(\tau)$  varies between -1 and 1. The base load  $u_{\text{el},k}^0$  can be changed by the reserves provider every 15 minutes,  $r_k$  is fixed at the beginning of the day, but time varying in 15 minute intervals, and  $w(\tau)$  updates every 2 seconds. As a result,  $u_{\text{el},k}(\tau)$  is also updated every 2 seconds.

### 5.2.2 System under consideration




Figure 5.1: Schematic of the system under consideration with heat pump, warm and cold water storage tanks, controlled and uncontrolled buildings with heat exchangers for heating and cooling, and a borehole field for seasonal storage.

The system under consideration, shown in Figure 5.1, is a small district heating and cooling system which is supplied by a central large-scale heat pump. On the evaporator side, the heat pump is connected to a cold thermal storage tank, and on the condenser

side to a warm thermal storage tank. To cover peak demands, the system is connected to an external heating and cooling grid via heat exchangers behind the condenser and evaporator respectively. Individual buildings and apartments draw water from either one of the storage tanks to serve their heating or cooling demand. We consider buildings that are controllable by the central system operator (yellow) and buildings that are uncontrollable<sup>2</sup> (blue), i.e. they are controlled individually by their own control system considered unknown to the system operator. Note, that we have only depicted one controlled and one uncontrolled building here for clarity, but there can be many. As the heating and cooling demand of the entire system rarely match each other, excess heating or cooling energy is stored in a borehole field for seasonal storage<sup>3</sup>.

### 5.2.3 Control scheme

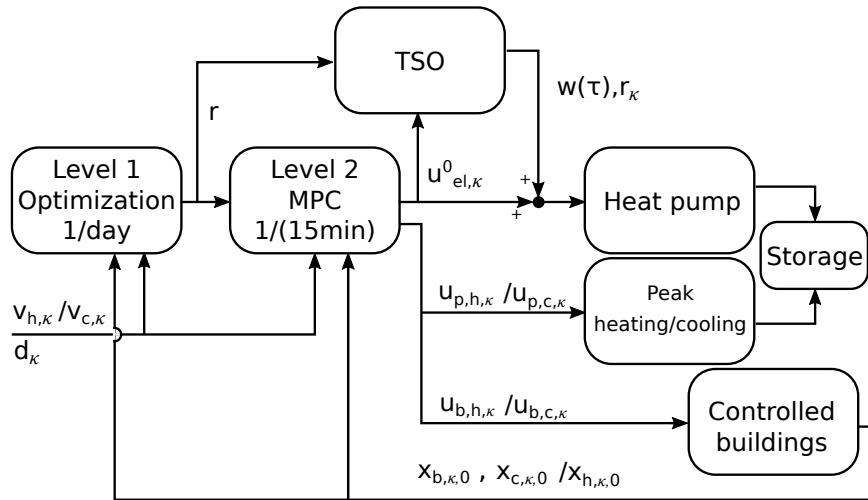


Figure 5.2: Hierarchical control scheme with two control Levels.

The presented two-level control scheme, as shown in Figure 5.2, is adapted from (Bünning et al., 2020d; Vrettos et al., 2018b) and Chapter 3. Here, we extend it to include the dynamics of buildings connected to the district heating/cooling system. Level 1 is a robust optimization problem which is solved at the beginning of the day with the purpose to communicate the reserves offered  $r$  for the following 24 hours to the TSO. It tries to maximize the reserves, while minimizing the costs for electricity and peak heating/cooling, maintaining sufficient temperature levels in the storage tanks to serve the uncontrollable buildings, and respecting comfort constraints of the controlled buildings. The optimization problem is formulated on the basis of models of the heat pump, storage

<sup>2</sup>Note that *controllable* and *uncontrollable* do not refer to the system-theoretic attributes here, but are rather a description of whether the central controller can actively control the considered buildings.

<sup>3</sup>There is a separate heat exchanger that decouples the shown water-based system and the borehole glycol cycle that is not shown in the Figure.

tanks, and controlled buildings. As inputs, it takes the initial states of the cold and warm storage tanks,  $x_{c,\kappa,0}$  and  $x_{h,\kappa,0}$  respectively, initial states of the controlled buildings  $x_{b,\kappa,0}$ , forecasts for the heating and cooling demand of the uncontrolled buildings,  $v_{h,\kappa}$  and  $v_{c,\kappa}$ , and forecasts of the disturbances to the controlled buildings  $d_\kappa$ . Note that, compared to  $k$ , introduced in (5.1), which describes the  $k$ -th element in a vector,  $\kappa$  denotes the time of optimization (which is also discretized in 15-minute time steps).

Level 2 is a receding horizon MPC controller based on the same optimization problem as Level 1 with the purpose of re-optimizing the heat pump base load ( $u_{el,\kappa}^0$ ), the peak heating and cooling ( $u_{p,h,\kappa}$  and  $u_{p,c,\kappa}$ ), and heating/cooling control inputs to the controlled buildings ( $u_{b,h,\kappa}$  and  $u_{b,c,\kappa}$ ) on the basis of updated disturbance forecasts and initial conditions. The reserves  $r$  of the current day are not decision variables in Level 2, as they have been fixed by Level 1.

For the heat pump, the peak heating and cooling heat exchangers, and the controlled buildings, we assume that lower-level controllers are in place, which track the desired set points with reasonable performance. Such controllers are discussed in Chapters 3 and 4, and are out of scope of this chapter.

## 5.3 Methodology

### 5.3.1 Modeling

For the optimization problems in the controller Levels 1 and 2, models of the considered system components are needed. These are discussed on the following.

#### Component models

For the storage tanks, we neglect thermal losses, because they happen on a much longer time scale and are small compared to charging and discharging by the heat pump and the buildings. Also the mass and heat capacity of the storage container itself are not modeled, because they are small compared to those of the water stored. The stored energy of the hot storage tank in continuous time can then be described in terms of its average temperature  $x_h(t)$  with

$$m_h c_{p,h} \frac{dx_h(t)}{dt} = u_{HP,h}(t) + u_{p,h}(t) - u_{b,h}(t) - v_{b,h}(t) + \delta_h(t), \quad (5.2)$$

where  $m_h$  and  $c_{p,h}$  denote the mass and specific heat capacity of the water,  $u_{HP,h}(t)$  and  $u_{p,h}(t)$  denote the heat supplied by the heat pump and the peak heat source respectively, and  $u_{b,h}(t)$  and  $v_{b,h}(t)$  are the sums of the heating demand of the controlled and uncontrolled buildings respectively. The uncertainty with respect to the heating demand of the

uncontrolled buildings is denoted by  $\delta_h(t)$ , which will be modeled as a box-constrained uncertainty set in the robust controller. Note that we have shown in Chapter 3, that modeling the average temperature is sufficient for this control task and a stratified tank model is not required. Moreover, errors induced by this simplification can be captured by  $\delta_h(t)$ . A model for the cold storage tank can be expressed accordingly with

$$m_c c_{p,c} \frac{dx_c(t)}{dt} = -u_{\text{HP},c}(t) - u_{p,c}(t) + u_{b,c}(t) + v_{b,c}(t) + \delta_c(t), \quad (5.3)$$

where the type of variables are identical to the ones in (5.2), but describe the cooling case. Note that all variables are considered to be positive and the direction of energy transfer is encoded in the signs.

The heating output of the heat pump can be described by the energy balance

$$u_{\text{HP},h}(t) = u_{\text{HP},c}(t) + u_{\text{HP},el}(t) + \tilde{e}(t) = \text{COP}_h u_{\text{HP},el}(t) + e(t), \quad (5.4)$$

where  $\text{COP}_h$  is the heating coefficient of performance,  $u_{\text{HP},el}(t)$  is the electrical power consumption, and  $u_{\text{HP},c}(t)$  is the cooling power. As discussed in Chapter 3, the COP is modeled as a constant. The error term  $e(t)$  can capture the resulting modeling error and potential additional thermal disturbances. It will also be modeled as a box-constrained uncertainty set in the controller. After rearranging (5.4), the cooling output of the heat pump can be described in terms of the electrical power by

$$u_{\text{HP},c}(t) = (\text{COP}_h - 1) u_{\text{HP},el}(t) + e(t). \quad (5.5)$$

To actively integrate the controlled buildings in the control scheme and to exploit their thermal inertia for reserves, models of their thermal dynamics are needed. The continuous time temperature of room  $j$  in a controlled building  $i$  can be described by

$$\begin{aligned} m_{b,i,j} c_{p,b,i,j} \frac{dx_{b,i,j}(t)}{dt} = & \alpha_{\text{amb},i,j} (d_{\text{amb}}(t) - x_{b,i,j}(t)) \\ & + \sum_{n \in \mathcal{N}_j} \alpha_{n,i,j} (x_{n,i,j}(t) - x_{b,i,j}(t)) \\ & + u_{b,h,i,j}(t) + u_{b,c,i,j}(t) \\ & + d_{\text{sol}}(t) + \gamma_{i,j}(t), \end{aligned} \quad (5.6)$$

where  $m_{b,i,j}$  and  $c_{p,b,i,j}$  denote mass and specific heat capacity of the room respectively,  $\alpha_{\text{amb},i,j}$  and  $d_{\text{amb}}$  are the heat transfer coefficient with the ambient and the ambient temperature respectively,  $n \in \mathcal{N}_j$  denotes the neighboring rooms of room  $j$ ,  $\alpha_{n,i,j}$  and  $x_{n,i,j}$

are the heat transfer coefficient and the temperature of a neighboring room respectively,  $u_{b,h,i,j}$  and  $u_{c,h,i,j}$  are heating and cooling control inputs,  $d_{\text{sol}}(t)$  are heat gains from solar irradiation and  $\gamma_{i,j}(t)$  is the model error.

### Whole system model

As all presented models are linear, they can be rearranged, discretized in time and written as a discrete time linear state space system of the form  $x_{k+1} = \tilde{A}x_k + \tilde{B}u_k$  (Borrelli et al., 2017). After stacking the discretized variables along the prediction horizon  $N$  in vectors  $\mathbb{R}^N$ , for example by defining the trajectory of the warm temperature storage as  $x_h := [x_{h,1}, \dots, x_{h,N}]^\top \in \mathbb{R}^N$  and the heating output of the heat pump as  $u_{\text{HP},h} := [u_{\text{HP},h,1}, \dots, u_{\text{HP},h,N}]^\top \in \mathbb{R}^N$ , we can describe the state evolution of all system components in terms of the initial state conditions, inputs and disturbances as

$$x_h = A_h x_{h,0} + B_h (u_{\text{HP},h} + u_{p,h} - u_{b,h} - v_{b,h} + \delta_h) \quad (5.7a)$$

$$u_{\text{HP},h} = (\text{COP}_h) u_{\text{HP},\text{el}} + e, \quad (5.7b)$$

$$x_c = A_c x_{c,0} + B_c (-u_{\text{HP},c} - u_{p,c} + u_{b,c} + v_{b,c} + \delta_c) \quad (5.7c)$$

$$u_{\text{HP},c} = (\text{COP}_h - 1) u_{\text{HP},\text{el}} + e, \quad (5.7d)$$

$$x_{b,i} = A_{b,i} x_{b,0} + B_{b,u,i} (u_{b,h,i} - u_{c,h,i}) + B_{b,d,i} d + \gamma_i, \quad \forall i \in \mathcal{B}, \quad (5.7e)$$

$$u_{b,h} = \sum_{i \in \mathcal{B}} u_{b,h,i}, \quad (5.7f)$$

$$u_{c,h} = \sum_{i \in \mathcal{B}} u_{c,h,i}, \quad (5.7g)$$

with the matrices  $A$  and  $B$  defined appropriately. Here, (5.7a) describes the state evolution of the warm temperature storage, with  $u_{b,h}$  denoting the sum of all heating inputs of the controlled buildings  $i$  in the set of buildings  $\mathcal{B}$ , as described in (5.7f). Accordingly, Equation (5.7c) describes the cold storage temperature evolution, with the cooling power of the controlled buildings described by (5.7g). Equations (5.7b) and (5.7d) describe the heating and cooling power of the heat pump. Equation (5.7e) denotes the temperature evolution in the controlled buildings, where  $x_{b,i}$  is a stacked vector of the room temperature trajectories of building  $i$ . Accordingly,  $u_{b,h,i}$  and  $u_{c,h,i}$  are stacked vectors of the heating and cooling inputs to the individual rooms, and  $d$  comprises the disturbances  $d_{\text{sol}}$  and  $d_{\text{amb}}$ .

### Model parameterization

As repeatedly discussed in this thesis, modeling buildings from first principles is tedious because each building is individual. However, in Chapter 4 we have demonstrated in experiments, that linear ARMAX models with physics-based inputs and constraints on the model coefficients are sufficient for predictive temperature control in buildings. These models can be generated directly from measurement data obtained under standard building operation, for example with hysteresis temperature control.

By treating the discretized version of (5.6) as an input-output model with one lag term, model output  $x_{b,i,j,k+1}$ , and model inputs  $x_{b,i,j,k}$ ,  $d_{amb,k}$ ,  $x_{n,i,j,k}$ ,  $u_{b,h,i,j,k}$ ,  $d_{sol,k}$ , we can identify the coefficients that multiply the model inputs with the method presented in Chapter 4 and construct the matrices  $A_{b,i}$ ,  $B_{b,u,i}$  and  $B_{b,d,i}$  for model (5.7e). Here, the disturbance  $d_{sol,k}$  comprises the one-hot encoding of the solar irradiation as described in Chapter 4.

As (5.2)-(5.5) (and (5.7a)-(5.7d)) describe commercially available equipment, they can be parameterized efficiently based on physics and manufacturer data.

### 5.3.2 Controller design

In the following, we will explain the design of the two-level controller and the integration of the models developed above.

#### Heating and cooling mode

The heating and cooling power of a heat pump are directly coupled through (5.4). The heat pump can, therefore, either be operated heating-driven or cooling-driven, which either leads to excess cooling power or excess heating power respectively. For the considered system, this means that the desired temperature constraints in the storage tanks cannot be guaranteed at the same time: if the connected buildings require significantly more heating than cooling, the temperature in the cold storage tank will drop as a consequence and vice versa. The excess cooling or heating energy is stored in the borehole field.<sup>4</sup>

For the presented controller, we therefore define a heating and a cooling mode: At the beginning of the day, when controller Level 1 is executed, we compute based on the heating and cooling demand forecasts  $v_{b,h}$  and  $v_{b,c}$ , and Equations (5.4) and (5.5),

---

<sup>4</sup>Note, that the borehole field is not specifically taken into account in this study for two reasons. First, under normal operating conditions of the considered demonstrator system, it is not considered either. A potential mismatch between charging and discharging is accepted at the cost of a sub-optimal borehole field temperature. Second, seasonal storage management requires different modeling techniques which are out of scope of this study. The interested reader is referred to (Flamm et al., 2021).

whether the sum of electrical energy for the day required to supply heating is larger than the one required to supply cooling. If yes, the controller is run in heating mode and the cold thermal storage is not considered in the controller (as we can then assume that it will always be sufficiently cold), if no, the controller is run in cooling mode.<sup>5</sup>

### Level 1 controller

The purpose of Level 1 is to optimize the offered reserves for the next 24 hours and to communicate them to the TSO. For this, we define an optimization problem that minimizes the cost of operation of the system while maintaining a sufficient temperature in the storage tank to meet the demand of the uncontrolled buildings, and while maintaining the comfort constraints in the controlled buildings. For the sake of space, we only discuss the heating case. The cooling case works accordingly by exchanging the dynamics for the heat pump and storage. The optimization problem is defined as

$$\begin{aligned} \min_{\substack{r, x_h, u_{\text{HP,el}}^0, \\ u_{\text{HP,h}}, \epsilon_h, z \\ x_{b,i}, u_{b,h,i}, u_{c,h,i}, \\ u_{b,h}, u_{c,h}, \epsilon_{b,i}}} \quad & f^{\text{el}\top} u_{\text{HP,el}}^0 - f^r \top r + f^{\text{p}\top} u_{\text{p,h}} + \lambda^{\text{h}\top} \epsilon_h + \lambda^{\text{b}\top} \epsilon_b \quad (5.8a) \\ \text{subject to} \quad & x_h = A_h x_{h,0} + B_h (u_{\text{HP,h}} + u_{\text{p,h}} - u_{b,h} - v_{b,h} + \delta_h), \quad (5.8b) \\ & u_{\text{HP,h}} = \text{COP}_h (u_{\text{HP,el}}^0 + \bar{w} \odot r) + e, \quad (5.8c) \\ & X_{h,\min} - \epsilon_h \leq x_h \leq X_{h,\max} + \epsilon_h, \quad (5.8d) \\ & z U_{\text{HP},\min} \leq u_{\text{HP,el}}^0 + w \odot r \leq z U_{\text{HP},\max}, \quad (5.8e) \\ & \tilde{z} R_{\min} \leq r \leq \tilde{z} R_{\max}, \quad (5.8f) \\ & z, \tilde{z} \in \mathbb{Z}_2^N, \quad (5.8g) \\ & \epsilon_h \geq 0, \quad (5.8h) \\ & x_{b,i} = A_{b,i} x_{b,0} + B_{b,u,i} (u_{b,h,i} - u_{c,h,i}) + B_{b,d,i} d + \gamma_i, \quad (5.8i) \\ & u_{b,h} = \sum_{i \in \mathcal{B}} u_{b,h,i}, \quad (5.8j) \\ & u_{c,h} = \sum_{i \in \mathcal{B}} u_{c,h,i}, \quad (5.8k) \\ & X_{b,i,\min} - \epsilon_{b,i} \leq x_{b,i} \leq X_{b,i,\max} + \epsilon_{b,i}, \quad (5.8l) \\ & \epsilon_{b,i} \geq 0, \quad (5.8m) \\ & U_{b,h,i,\min} \leq u_{b,h,i} \leq U_{b,h,i,\max}, \quad (5.8n) \\ & U_{b,c,i,\min} \leq u_{b,c,i} \leq U_{b,c,i,\max}, \quad (5.8o) \end{aligned}$$

---

<sup>5</sup>A common strategy in practical applications is to set heating or cooling modes based on the ambient temperature. However, this is less exact than our method.



$$\forall i \in \mathcal{B}, \quad (5.8p)$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta_h \in \Delta_h, \forall e \in E, \forall \gamma_i \in \Gamma_i. \quad (5.8q)$$

Here,  $f^{\text{el}}$ ,  $f^{\text{p}}$ ,  $\lambda^{\text{h}}$  and  $\lambda^{\text{b}}$  denote the costs for electricity consumption, peak heating (through the external heating grid), temperature constraint violations in the storage tank, and temperature constraint violations in the controlled buildings respectively. The benefits for reserves offered are denoted by  $f^r$ .

Equations (5.8b)-(5.8h) describe the constraints directly related to the heat pump and water storage. Equation (5.8b) defines the dynamics of the storage, as described in (5.7a), and (5.8c) describes the dynamics of the heat pump with the reserve tracking signal (5.1) inserted for  $u_{\text{HP,el}}$ . The symbol  $\odot$  denotes the operator for element-wise multiplication. Note that the regulation signal  $w(\tau)$  is updated every 2 seconds and takes values in the interval  $[-1, 1]$ . However, the decision variables in (5.8) refer to quantities that are updated every 15 minutes. Therefore, to meet the robust temperature constraint of the storage tank (5.8d), the average value of  $w(\tau)$  over a 15 minute interval, denoted by  $\bar{w}$ , is used in (5.8c) because it is more descriptive than the instantaneous value. The temperature limits of the storage tank  $X_{\text{w,min}}$  and  $X_{\text{w,max}}$  in (5.8d) are defined by the lowest possible operating temperature to supply heating to the connected buildings, and the highest desirable supply temperature of the heat pump. The slack variable  $\epsilon_{\text{h}}$  ensures feasibility of this constraint. Constraint (5.8e) defines the operational limits of the heat pump, with  $U_{\text{HP,min}}$  and  $U_{\text{HP,max}}$  being the lower and higher capacity limit respectively. Here,  $w$  is used instead of  $\bar{w}$ , because the operational constraint has to be maintained for all realizations of the regulation signal, not just for the average. The binary variable  $z$  is introduced to allow an on-off condition while also having the lower capacity limit. As close-to zero offered reserves can lead to bad tracking performance (Bünning et al., 2020c), (5.8f) also implements limits on the offered reserves, with  $R_{\text{min}}$  and  $R_{\text{max}}$  denoting the lower and upper limit respectively, and  $\tilde{z}$  being a binary that allows  $r$  to attain zero.

Equations (5.8i)-(5.8p) describe constraints and qualifiers directly related to the controlled buildings. Constraints (5.8i)-(5.8k) define the dynamics of the buildings as defined in (5.7e)-(5.7g). The comfort constraints are defined in (5.8l) with  $X_{\text{b},i,\text{min}}$  and  $X_{\text{b},i,\text{max}}$  denoting lower and upper comfort constraints respectively and  $\epsilon_{\text{b},i}$  being a slack variable to ensure feasibility. Input constraints for heating and cooling energy are defined in (5.8n)-(5.8o), with  $U_{\text{b,h},i,\text{min}}$  and  $U_{\text{b,h},i,\text{max}}$  denoting lower and upper bounds for the heating inputs respectively, and  $U_{\text{b,c},i,\text{min}}$  and  $U_{\text{b,c},i,\text{max}}$  denoting the limits for cooling.

The presented constraints have to hold for all uncertainty realizations, introduced by the regulation signal  $w$ , thermal disturbances in the heat pump  $e$ , uncertainty in the demand prediction  $\delta_{\text{h}}$  and modeling error of the controlled buildings  $\gamma_i$ , as defined in the qualifier (5.8q).

As discussed in Chapter 3, one drawback of the method is that, the effect of the uncertainties on the states  $x_h$  and  $x_{b,i}$  compound along the prediction horizon through the integrators implicit in (5.8b) and (5.8i). As explained in Chapter 3, here, we therefore also introduce affine policies on uncertain variables in the form of matrices  $D$ , which change the optimization problem to

$$\min_{\substack{r, x_h, u_{\text{HP,el}}^0, \\ u_{\text{HP,h}}, \epsilon_h, z \\ x_{b,i}, u_{b,h,i}, u_{c,h,i}, \\ u_{b,h}, u_{c,h}, \epsilon_{b,i}, D}} f^{\text{el}\top} u_{\text{HP,el}}^0 - f^r \top r + f^p \top u_{p,h} + \lambda^h \top \epsilon_h + \lambda^b \top \epsilon_b \quad (5.9a)$$

$$\text{subject to } x_h = A_h x_{h,0} + B_h (u_{\text{HP,h}} + u_{p,h} - u_{b,h} - v_{b,h} + \delta_h), \quad (5.9b)$$

$$u_{\text{HP,h}} = \text{COP}_h (u_{\text{HP,el}}^0 + \bar{w} \odot r + D_{w,\text{HP}} \bar{w} + D_{\delta_h e, \text{HP}} (\delta_h + e) + \sum_{i \in \mathcal{B}} D_{\gamma_i, \text{HP}} \gamma_i) + e, \quad (5.9c)$$

$$X_{h,\min} - \epsilon_h \leq x_h \leq X_{h,\max} + \epsilon_h, \quad (5.9d)$$

$$z U_{\text{HP},\min} \leq u_{\text{HP,el}}^0 + w \odot r + D_{w,\text{HP}} \bar{w} + D_{\delta_h e, \text{HP}} (\delta_h + e) + \sum_{i \in \mathcal{B}} D_{\gamma_i, \text{HP}} \gamma_i \leq z U_{\text{HP},\max}, \quad (5.9e)$$

$$\tilde{z} R_{\min} \leq r \leq \tilde{z} R_{\max}, \quad (5.9f)$$

$$z, \tilde{z} \in \mathbb{Z}_2^N, \quad (5.9g)$$

$$\epsilon_h \geq 0, \quad (5.9h)$$

$$x_{b,i} = A_{b,i} x_{b,0} + B_{b,u,i} (u_{b,h,i} - u_{c,h,i} + (D_{w,b,h,i} - D_{w,b,c,i}) \bar{w} + (D_{\delta_h e, b, h, i} - D_{\delta_h e, b, c, i}) (\delta_h + e) + (D_{\gamma, b, h, i} - D_{\gamma, b, c, i}) \gamma_i) + B_{b,d,i} d + \gamma_i, \quad (5.9i)$$

$$u_{b,h} = \sum_{i \in \mathcal{B}} (u_{b,h,i} + D_{w,b,h,i} \bar{w} + D_{\delta_h e, b, h, i} (\delta_h + e) + D_{\gamma, b, h, i} \gamma_i), \quad (5.9j)$$

$$u_{b,c} = \sum_{i \in \mathcal{B}} (u_{b,c,i} + D_{w,b,c,i} \bar{w} + D_{\delta_h e, b, c, i} (\delta_h + e) + D_{\gamma, b, c, i} \gamma_i), \quad (5.9k)$$

$$X_{b,i,\min} - \epsilon_{b,i} \leq x_{b,i} \leq X_{b,i,\max} + \epsilon_{b,i}, \quad (5.9l)$$

$$\epsilon_{b,i} \geq 0, \quad (5.9m)$$

$$U_{b,h,i,\min} \leq u_{b,h,i} + D_{w,b,h,i} \bar{w} + D_{\delta_h e, b, h, i} (\delta_h + e) + D_{\gamma, b, h, i} \gamma_i \leq U_{b,h,i,\max}, \quad (5.9n)$$

$$U_{b,c,i,\min} \leq u_{b,c,i} + D_{w,b,c,i} \bar{w} + D_{\delta_h e, b, c, i} (\delta_h + e) + D_{\gamma, b, c, i} \gamma_i \leq U_{b,c,i,\max}, \quad (5.9o)$$

$$\forall i \in \mathcal{B}, \quad (5.9p)$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta_h \in \Delta_h, \forall e \in E, \forall \gamma_i \in \Gamma_i. \quad (5.9q)$$

In (5.9a),  $D$  denotes all affine policies defined in the constraints (5.9b)-(5.9o). As discussed in Chapter 3, the  $D$  matrices are constrained to be causal, i.e. strictly lower triangular, such that the controller can only implement feedback policies on past uncertainty realizations with respect to the considered time step, but not to future ones. In addition to the dynamics of the storage tank and the temperature of the controlled buildings, also the capacity constraints for the heat pump and the input constraints of the buildings have to hold for the introduced policies. Note, that in principle, each controlled building could also react to the disturbance realization observed in the other buildings. For simplicity of notation we are neglecting this possibility and just show policies where each building only reacts to their own disturbance realizations.

As in Chapter 3, the uncertainty sets  $W$ ,  $\bar{W}$ ,  $\Delta_h$ ,  $E$ , and  $\Gamma_i$  generally allow any convex sets. For box-constrained sets, problem (5.9) can be reformulated as a Mixed Integer Linear Program. Many modern optimization tools (Löfberg, 2004; Goh and Sim, 2011) perform this reformulation automatically, for example through *explicit maximization* Löfberg (2012).

The final optimization scheme used for controller Level 1 is (5.9). (Problem (5.8) was presented to ease following the modeling process.)

## Level 2 controller

Level 2 is a receding horizon MPC scheme. Its purpose is to update the control inputs of the controlled buildings, as well as the heat pump base load  $u_{\text{HP,el}}^0$  on the basis of new disturbance forecasts, which are updated every 15 minutes, and new measurements of initial states (which implicitly also encode the uncertainty realizations). Level 2 uses the same optimization problem as Level 1, i.e. (5.9). However, the offered reserves  $r$  have been fixed by Level 1 for the whole day. This means that at optimization time  $\kappa$ , the first  $N - \kappa$  elements of  $r$  are constrained to the values decided on by Level 1. The last  $\kappa$  elements can be chosen by Level 2, or can be constrained to zero for example for the sake of computational efficiency.

Note, that the scheme also works with a shrinking horizon in Level 2, as was shown in Chapter 3. Compared to this, a receding horizon scheme where the not-fixed reserve variables are padded with zeros has the advantage of not depleting the storage towards the end of the day when the shrinking horizon becomes small. This is at the expense of a slightly more complex<sup>6</sup> optimization problem. However, the experiments in Chapter 3

<sup>6</sup>On the other hand, for the receding horizon approach the optimization problem compiling process is considerably faster as only one optimization problem needs to be constructed, whereas  $N$  problems need to be built for the shrinking horizon approach. As deriving the robust counterpart of Problem (5.9) is computationally expensive, there is a trade-off between both approaches.



Figure 5.3: DFAB HOUSE at NEST at Empa in Switzerland. Copyright: Roman Keller.

and this chapter show that both approaches are viable in practice.

## 5.4 Numerical case study and results

In Chapter 3, we have presented a reserve provision scheme which offers reserves in large building energy systems or district heating systems, purely by exploiting the thermal flexibility of buffer storage tanks. In this case study, we therefore investigate the added potential from exploiting the thermal flexibility of a subset of the connected buildings as presented in Section 5.3.

### 5.4.1 Configuration

The case study is conducted on the basis of historical data of one year (2019-06-24 to 2020-06-24) from the NEST demonstrator (Figure 2.2). We use the same heating system configuration with the heat pump and warm storage tanks as described in Chapter 3. However, in contrast to Chapter 3, we also consider the cooling system, which also comprises two 1100 litre Matica water tanks, connected to the evaporator side of the heat pump. The system resembles the configuration described in Figure 5.1.

As controlled buildings, we consider the units UMAR (Urban Mining and Recycling)(Heisel and Rau-Oberhuber, 2020), depicted on the second floor of Figure 2.2, and DFAB (Digital Fabrication and Living)(Graser et al., 2021), depicted in Figure 5.3. As described in Chapter 4, UMAR is a two-bedroom apartment with a joint living and kitchen area, a hallway and two bathrooms. Its heating and cooling system comprises ceiling heating/cooling panels which behave similarly to conventional wall radiators. DFAB is a three-story apartment comprising four bedrooms, a living and kitchen area

Table 5.1: Parameters for controller Level 1 related to heat pump and storage.

$N = 96,$	$\text{COP}_h = 3.53,$	$\tilde{A}_h, \tilde{A}_c = 1,$
$\lambda^h = 20,$	$\text{COP}_c = 2.73,$	$\tilde{B}_h, \tilde{B}_c = 0.0978 \frac{K}{kW},$
$f^{\text{el}} = 1,$	$W = [-1, 1],$	$X_{h,\text{min}} = 28^\circ C,$
$f^r = 1.5,$	$\bar{W} = [-0.25, 0.25],$	$X_{h,\text{max}} = 38^\circ C,$
$f^p = 1,$	$U_{\text{min}} = 8.2kW,$	$X_{c,\text{min}} = 7^\circ C,$
$R_{\text{min}} = 0.4kW,$	$U_{\text{max}} = 12.8kW,$	$X_{c,\text{max}} = 18^\circ C,$
$x_{h,0} = 33^\circ C$	$x_{c,0} = 12.5^\circ C$	$E \oplus \Delta_{h,c} = [-4.0, 4.0] kW$

and four bathrooms. It is heated and cooled by a floor heating/cooling system. All other units of NEST, i.e. several residential apartments, office units, and a fitness center, are considered uncontrollable.

In the numerical study, we only consider controller Level 1, as we intend to investigate the reserves offered. The parameters for the controller are shown in Table 5.1. The cost-function related parameters  $f^{\text{el}}$ ,  $f^r$  and  $f^p$  were chosen on the basis of the studies in Chapter 3 to maintain comparability. They are not directly related to a specific local energy product, but rather reflect the assumption that the benefits for offering reserves are higher than the base price for electricity, and back-up heating/cooling being more expensive (after considering the COP of the heat pump) than using the heat pump. The penalty for using the storage slack variable is set to 20, such that it is only used in case of inevitable infeasibility otherwise. The values for  $\text{COP}_h$ ,  $\text{COP}_c$ ,  $U_{\text{min}}$ ,  $U_{\text{max}}$  and  $R_{\text{min}}$  were set based on preliminary heat pump experiments. Note that the simplified relation  $\text{COP}_c = \text{COP}_h - 1$  does not hold here due to thermal losses in the heat pump. The bounds for  $W$  are defined by the regulation signal RegD by PJM and the averaged uncertainty set  $\bar{W}$  is set by analyzing historical regulation signals, as described in (Vrettos et al., 2018a). The storage properties  $\tilde{A}_h, \tilde{A}_c$  follow from the assumption of no thermal losses, and  $\tilde{B}_h, \tilde{B}_c$  are set by calculating the heat capacity based on the specific heat capacity of water and the tank volumes. The temperature constraints of the storage tanks  $X_{h,\text{min}}, X_{h,\text{max}}, X_{c,\text{min}}, X_{c,\text{max}}$  are based on the operational constraints used in the building to guarantee sufficient heating and cooling temperatures. The initial conditions of the storage tanks are set to  $33^\circ C$  and  $12.5^\circ C$ , for the heating and cooling case respectively, which is in the middle of their temperature range. The Minkowski sum  $E \oplus \Delta_{h,c}$ , which we use here because  $e$  and  $\delta_{h/c}$  appear in the same way in (5.9), are based on the studies in Chapter 3, and are a result of analyzing historical prediction errors of the heating demand forecasts of Chapter 2, and measurements of the thermal output of the heat pump. To increase the reliability of the heat pump, we also introduce an additional constraint that requires  $z_k$  to be constant during each thirty minute interval. The lower reserve constraint  $R_{\text{min}}$  is set to  $0.4kW$  as lower reserves lead to large relative tracking errors in preliminary experiments (Bünning et al., 2020c).

The models for UMAR and DFAB are parameterized based on historical measurement

data and the method of physics-informed linear regression as presented in Chapter 4, which has been validated in various experiments (Bünning et al., 2021a; Lefebure et al., 2021). From these experiments, we also know that a deterministic MPC formulation is suitable to maintain comfort in the units. The uncertainties related directly to the controlled buildings,  $\gamma_i$ , are therefore set to zero; consequently also the related affine policies disappear in (5.9). The heating and cooling input constraints are determined on the basis of preliminary experiments and vary for each room and each building unit. The initial temperature is set to 22.5°C in all rooms, and the comfort constraints are 20°C - 25°C. The constraint violation penalty  $\lambda^b$  is set to 20 for all rooms in both buildings. All disturbances, i.e. the heating demand of the uncontrolled buildings, ambient temperature and solar irradiation are historical measurements from the building site.

For the implementation of the optimization problem, we use Matlab<sup>®</sup> with YALMIP Löfberg (2012), which automatically derives robust counterparts. As all uncertain variables are box-constrained, the robust counterparts become MILPs, which we solve with Gurobi<sup>®</sup> 9.0.3. The solver stops the optimization if the gap to the lower MIP bound is less than 1% or if a time limit of 15 minutes is reached, in which case the best feasible solution is taken as the result. To increase computational efficiency, the affine policies are set to be equal between the rooms of each individual controlled building. Moreover, the number of feedback steps in the policies is limited to 4 in the storage and 2 in the units.

## 5.4.2 Results and discussion

In the case study, we compare the case where UMAR and DFAB are controllable buildings in (5.9) to the case where the entire NEST building including UMAR and DFAB is uncontrollable (as described in Chapter 3).

Figure 5.4 shows the results of the study. Subplot (a), depicts the mean daily heating and cooling demand of the entire NEST building with UMAR and DFAB included. As can be expected for a building in central Europe, the total heating demand is higher than the cooling demand. The orange shading indicates whether Level 1 was set to heating mode (orange) or cooling mode (no shading).

From subplot (b) it can be seen that the cost function<sup>7</sup> is generally slightly lower for the case where DFAB and UMAR are controlled (shown in green) to the case where the entire building is uncontrolled (dashed black). The total cost reduction over the course of the entire year is 8%. There are exceptions where the case with controlled units is more expensive, for example on several days after 2020-04. In theory, the unit-controlled

---

<sup>7</sup>We do not consider the constraint violation terms here as they are different for the case with controlled buildings and without controlled buildings. Moreover, in the case where the solver does not find the global optimal solution, small constraint violations lead to large distortions in the cost function, which makes comparison difficult.




Figure 5.4: Results of the numerical case study. (a): Average daily heating demand of entire NEST in orange, average daily cooling demand in dashed blue. Orange shaded areas indicate that Level 1 runs in heating mode, otherwise cooling mode. (b): Cost function without constraint violations taken into account for Level 1 considering DFAB and UMAR as controllable (green) and uncontrollable (dashed black). (c): Daily sum of offered reserves by Level 1 considering DFAB and UMAR as controllable (green) and uncontrollable (dashed black). (d): Base daily electricity consumption of the heat pump as planned by Level 1 considering DFAB and UMAR as controllable (green) and uncontrollable (dashed black). (e): Peak heating and cooling consumption as planned by Level 1 considering DFAB and UMAR as controllable (green) and uncontrollable (dashed black).

case should always be less or at least equally expensive, because Level 1 should be able to replicate the control action of the baseline hysteresis controllers implemented in UMAR and DFAB. This solution should lead to the same cost as the uncontrolled case under the assumption that there is no significant model mismatch.<sup>8</sup> The higher costs could be a result of the solver not finding the optimal solution within the time limit and using the peak heater/cooler too much (see subplot (e)). Moreover, in some instances<sup>9</sup> the Level 1 controller also optimizing DFAB and UMAR could not find any solution. In this case, we set the solution to be equal to the one obtained by Level 1 only optimizing the storage tanks. We will further discuss this issue in Section 5.6.

As can be seen in subplot (c), the reserves offered are in total 2.7 times higher when UMAR and DFAB are actively controlled by the reserve scheduling scheme. This is remarkable as UMAR and DFAB only account for 7% of the total heating demand and 12% of the total cooling demand of NEST. The ratio of the offered reserves to the base electricity consumption, shown in subplot (d), is increased from 2.8% in the case of no unit control to 7.3% in the case where UMAR and DFAB are controlled. This increase is not limited to specific periods but can be observed over the course of the whole year. In the period after 2020-05 with low heating and cooling demands, the reserves offered without unit control are close to zero, while in the case where UMAR and DFAB are controlled, reserves are still being offered.

The total electricity consumption is 0.3% higher for the case where the units are controlled by Level 1, which indicates that the heat pump does not need to be run at considerably higher set points to offer the reserves. The backup heat exchangers are used 16% less. This advantage mainly stems from the time period between 2019-12 and 2020-02, where the heating demand of all connected units is high. This result shows that exploiting the thermal flexibility of the controlled buildings not only helps increasing the offered electrical reserves, but can also decrease the need for peak power, which is intuitive because the units can shift their demand to avoid peaks. This observation is also interesting for the emerging research area of thermal reserves in district heating systems supplied by volatile thermal renewable sources (Ilić, 2020; Vandermeulen et al., 2018).

## 5.5 Experimental case study and results

To demonstrate the viability of the control approach and to investigate the controller behaviour, we have conducted an experiment in the heating case on the same demonstrator system as described in Section 5.4.

---

<sup>8</sup>This should be the case as the models have been validated in experiment repeatedly (Bünning et al., 2021a; Lefebure et al., 2021).

<sup>9</sup>In particular in 13% of the days, mostly occurring at very low heating and cooling demands.



### 5.5.1 Configuration

The configuration generally follows the one presented in the previous Section 5.4, with the exception that only DFAB is considered as a controllable building and the rest of the NEST demonstrator is treated as uncontrollable. Moreover, also the peak heating and cooling heat exchangers are not used<sup>10</sup> and the lower room temperature limit is set to 21°C. To solve Levels 1 and 2 in real time, forecasts of the heating demand of the uncontrolled building and the ambient conditions are necessary. The ambient conditions are forecast by the national weather service MeteoSwiss and are updated every twelve hours. The heating demand is forecast with the online-corrected Artificial Neural Network (ANN) method presented in Chapter 2, which is updated every 15 minutes. The method uses as inputs the forecast ambient temperature, the hour of the day (one hot encoded), the measured heating demand one day ago at the same time, the measured heating demand one week ago at the same time, and a binary variable that indicates whether it is a working or a non-working day. As a regulation signal, the RegD signal by PJM from the 27th of January 2019 is used.<sup>11</sup>

Level 1 is executed 10 minutes before the beginning of the day, while Level 2 is repeatedly executed every 15 minutes, 5 minutes before the control inputs are implemented on the system (time  $\kappa$ ). Different optimization times are used as Level 1 is computationally more heavy, because  $r$  is not fixed yet and as a consequence all binary variables are open. Moreover, compared to Level 1, the prediction horizon of Level 2 is reduced from 24 hours to 12 hours.<sup>12</sup> The best feasible solution found by the solver within the time limit is implemented on the system. To estimate the initial states of the storage and the controlled buildings, their states are measured at  $\kappa$ -15 minutes and  $\kappa$ -10 minutes for Level 1, and  $\kappa$ -15 and  $\kappa$ -5 minutes for Level 2, and linearly extrapolated to time  $\kappa$ .

The optimization problems Level 1 and 2 are implemented in Matlab, as described in Section 5.4, while the lower level controllers handling the communication with NEST are implemented in Python 3 through an OPC-UA client (Leitner and Mahnke, 2006). Data exchange between Matlab and Python is facilitated through shared csv files.

### 5.5.2 Results and discussion

Figure 5.5 shows the results of the experiment. From the heating demand, depicted in dashed blue in Figure 5.5 (a), it can be seen that the experiment covers a wide range of demands; from less than 20 kW, which is below the lower thermal capacity limit of

<sup>10</sup>In practice, these serve as backups operated by the fallback system, for the case when experiments go wrong.

<sup>11</sup>As in Chapter 3, this choice is made for convenience, as the signal starts and ends with a value of 1 and is thus continuous when repeated.

<sup>12</sup>This has little effect on the optimality of the implemented control inputs, as the time constant of both storage and units is considerably smaller than 12 hours.




Figure 5.5: Results of the control experiment at NEST. (a): actual heating demand of the building in dashed blue, initial forecast heating demand in Level 1 in orange, updated forecasts every 15 minutes for Level 2 in grey, specifically mentioned forecast in red. (b). Offered reserves in blue. (c): Average storage temperature in black, temperature constraints in dotted black, temperatures of individual storage layers in transparent colors. (d): Electrical power of heat pump in dashed blue, set point to be tracked in orange, range of possible set points as a result of the reserves offered in dotted red. (e): Comfort constraints for DFAB rooms in dotted black, room temperatures in different colors. (f): Total heating demand of DFAB in blue.




Figure 5.6: Room temperature examples from the units not directly controlled by the controller. The office unit “Meet2Create” is shown in blue, a permanent meeting room in orange, the residential unit “UMAR” in green, the office unit “Solace” in red, and the residential unit “Vision Wood” in purple.




Figure 5.7: Detailed plot of the heat pump following the regulation signal. Electrical power of heat pump in dashed blue, set point to be tracked in orange, range of possible set points as a result of the reserves offered in dotted red.

the heat pump, to more than 50kW, which is above the upper capacity limit of the heat pump (as we only consider the first compressor stage here). As a consequence, also the reserves offered, depicted in 5.5 (b) vary significantly. In the first third of the day, when the initial demand prediction (orange) is at an intermediate level, many reserves are offered, while during the last two thirds of the day, when the demand prediction is constantly close to the upper capacity limit of the heat pump, very little reserves are offered.

Subplot (c) shows the average warm storage temperature in solid black, and the temperature measurements of the individual layers in transparent colors. It can be seen that the temperature constraints, depicted in dotted black, are maintained during most times, except during the periods marked with cyan and orange backgrounds. In the period marked in cyan, the effect of the demand prediction error being outside the bounds of the uncertainty set can be noticed. At the time of optimization, the heating demand of the building was low and as a consequence also the demand prediction, shown in red in subplot (a), was low, because it is corrected based on the last measured

prediction error<sup>13</sup> (see Chapter 2). Thus, Level 2 switches off the heat pump, whose electricity consumption is shown in dashed blue in subplot (d), as no reserves were offered in the following 15-minute interval. However, shortly after, the heating demand of NEST increases by a larger amount than covered by the uncertainty set  $E \oplus \Delta_h$ , causing the storage temperature to drop below the constraints. Subplot (e), which shows the temperature trajectories of the rooms in DFAB demonstrates that this short constraint violation in the storage tank has no negative effect on the comfort in the connected buildings. During the entire experiment, the room temperatures are kept between the comfort constraints. This is also the case for the other, uncontrolled, apartments and office units connected to the system, as shown for some examples of room temperatures in Figure 5.6. The local controllers (which are considered unknown to Level 1 and 2) are able to keep the temperatures between the comfort constraints because the warm storage temperature is sufficiently high. Small violations of the temperature constraint in the storage do not cause comfort constraint violations in the connected buildings because of their thermal inertia.

In the period marked in orange in Figure 5.5, it can be seen that the Level 2 controller modulates the energy consumption in DFAB to maintain the temperature constraints in the storage: The heat pump can not maintain the storage temperature because the heating demand of NEST exceeds its thermal capacity limit. As a consequence, the heating energy consumption of DFAB is set to zero, as can be seen in subplot (f). Another example for this modulation of the controllable building heat consumption depending on the heating demand of the other NEST units and the storage temperature, can also be seen in the purple period. Before midnight, the heating demand of NEST is high and the water storage is depleting although the heat pump runs at full capacity. As a consequence, the energy consumption in DFAB is kept low by the Level 2 controller. As soon as the heating demand of NEST decreases, and the storage temperature is rising, the controller increases the heating consumption of DFAB again.

Generally, the controllers keep the room temperatures in DFAB close to the lower constraint to save energy with the exception of the room with the smallest heat capacity, which is depicted in orange in subplot (e). This can be explained by the affine policies being constrained to be identical for all rooms. To maximize recourse on uncertainties (i.e. to make the coefficients in the  $D$  matrices large), the room with a low thermal capacity has to be kept far from the comfort constraints (because reacting to uncertainties will drive the temperature towards the constraints fast), while the rooms with a high thermal capacity can stay close to the lower bound. This issue could potentially be mitigated by scaling the the policies with the thermal capacity of the rooms. Other

---

<sup>13</sup>Note, that there is at times a difference between the measured heating demand shown in dashed blue in Figure 5.5 (a) and the measured heating demand seen by the forecasting method. The depicted heating demand is averaged over a 15-minute interval, while the forecasting method uses instantaneous demand measurements. This choice needs to be reconsidered in future research.

possible explanations include that the room is directed towards the sun, resulting in higher heating gains compared to the other rooms, and that the identified model indicates good thermal insulation towards the ambient and a high heat transfer coefficient towards neighboring rooms. This could cause the optimizer to run the room at a higher temperature to implicitly heat the neighboring rooms.

Subplot (d) shows the heat pump following the regulation signal during times where reserves are offered (see subplot (b)). As shown in detail for the first 1.5 hours of the experiment in Figure 5.7, this reserve tracking happens with high accuracy by manipulating the compressor speed in the heat pump on a lower controller level. The reserves offered together with the base load of the heat pump, determine the range of operation, which is limited by the dotted red line. The tracking signal for the heat pump, as defined by (5.1) is shown in orange and the actual electricity consumption of the heat pump in dashed blue. The lower-level controller and the tracking performance have been discussed in detail in Chapter 3 and are not discussed here to avoid repetition.

## 5.6 Limitations

In Section 5.4, we have shown the potential of exploiting the thermal inertia of district heating/cooling-connected buildings to offer electrical reserves with the example of a real system. In Section 5.5, we have demonstrated in an experiment on the same system, in which a wide range of heating demand situations are covered, the control-performance of the presented reserve scheduling and control scheme. While longer experiments would further benefit the confidence in the scheme, together with previous experimental studies performed on the operation of the heat pump (Bünning et al. (2020d), Chapter 3), on predictive control of the units (Bünning et al. (2021a), Chapter 4), on the demand forecasting approach (Bünning et al. (2020a), Chapter 2), and on a combined distributed operation of all entities (Lefebure et al., 2021), the presented experiments promise the viability of the approach for real application.

However, the computational complexity of the reserve scheduling problem needs to be addressed in future research. For the presented numerical case study, sometimes a feasible solution could not be found within 15 minutes when the predicted heating/cooling demand is very low, although only two buildings were considered to be controllable and the affine policies were constrained to be the same for the rooms of each building. Although the optimization time could be further increased in practice, as Level 1 only has to be solved once per day (and Level 2 is less computationally heavy due to many fixed binary variables and the option of a shortened prediction horizon), the presented approach is likely limited to small community heating/cooling systems. For larger scale applications, distributed optimization approaches could be investigated, which besides sharing the computational burden between agents, would also suit the distributed struc-

ture of the district heating/and cooling system. Implementing the optimization problem directly in a solver language instead of YALMIP could further improve optimization time. Alternatively, the reserve provision could be disabled for days with particularly low heating/cooling demands, where the solver struggles. Here, the expected reserves offered are anyways low, because the heat pump has to operate at the lower capacity limit.

## **5.7 Conclusion**

In this Chapter, we have combined the methods of the previous chapters and presented a two-level Robust MPC approach with affine policies to offer frequency reserves with a district heating/cooling system supplied by a heat pump. The approach exploits the thermal inertia of buffer storage tanks and a subset of the connected buildings, which are controlled by the central controller directly. In a numerical case study based on historical data of a real system, we have shown that the approach significantly increases the offered reserves, compared to a case where only the thermal inertia of buffer storage tanks is exploited. Moreover, also the operational costs are decreased. In an experiment on the same system, we have demonstrated that the controller indeed offers and delivers electrical reserves, while maintaining comfort in all of the buildings connected to the district system. Future research should focus on the scalability of the approach, potentially investigating distributed optimization schemes.

## Conclusions and future directions

### 6.1 Summary and conclusion

In this thesis, we have developed and compared data-driven methods for building demand forecasting and building thermal dynamics modeling for predictive control. By combining physics-based modeling of industrial components with data-driven modeling of buildings and advanced control concepts, we have reduced the manual modeling effort considerably, leading to controllers that are viable for practical application. Addressing the objectives of reducing building energy consumption and providing electrical reserves, the presented methods were validated in experiments in building energy and district energy applications.

In Chapter 2, we presented two online correction methods, one based on error-autocorrelation and one based on online learning, for ANN-based heating demand forecasts. The methods were validated with measurement data of a variety of buildings, and it was found that they significantly increase prediction accuracy and reduce the variance of accuracy between different networks trained on the same data set. The methods also significantly reduce the amount of training data needed compared to uncorrected forecasts.

In Chapter 3, we combined these demand forecasts with Robust MPC to offer electrical reserves with a system comprising a heat pump and buffer storage tanks without modeling the dynamics of the connected buildings. The approach was validated in an experiment, where it was shown that the temperature constraints in the storage tanks are respected, ensuring the comfort in the connected buildings, while reserves are being offered and a regulation signal is tracked successfully by the heat pump.

In Chapter 4, we presented Machine Learning methods based on Random Forests and Input Convex Neural Networks for data-driven model generation for building MPC, as well as a method based on physics-informed ARMAX. In extensive heating and cooling experiments with predictive building control and time-varying comfort constraints, all methods maintained occupant comfort and at the same time significantly reduced the heating and cooling energy consumption compared to a baseline hysteresis controller.

However, the physics-informed ARMAX model significantly outperforms the other methods in terms of sample efficiency.

In Chapter 5, we combined the methods of the previous three chapters to offer electrical reserves with an electrified district heating/cooling system, exploiting the thermal inertia of warm and cold water storage and a subset of the connected buildings. The buildings were modeled with data-driven methods, whereas the heat pump and tanks were modeled with first principles. The potential to offer electrical reserves was investigated in a numerical case study on the basis of historical measurement data of a small district heating system. It was found that by exploiting the thermal inertia of a subset of the buildings, the reserves offered could be significantly increased compared to Chapter 3. Moreover, the approach was validated in experiment.

In conclusion, we have made the following contributions. We have demonstrated that data-driven methods can lead to viable predictive controllers for building thermal control that outperform state-of-industry controllers. We have also seen that they are transferable: The best-performing modeling method for predictive control, physics-informed ARMAX, was validated in two different apartment units, UMAR (Chapter 4) and DFAB (Chapter 5), in both of which it performed well. For the industrial application of data-driven MPC, it is especially encouraging that no excitation experiments with the purpose to create informative training data were conducted. All training data used was generated under standard building operation, i.e. with hysteresis-based controllers. The reasonable control performance of the other two methods, RF and ICNN, given their lower prediction accuracy, also shows that a perfect model is not necessary for predictive building thermal management. This observation traces its origins to the fact that building dynamics are inherently stable, i.e. *wrong* control inputs do not take the system state to unsafe trajectories. Moreover, as the dynamics are slow, the required sampling times are long, and the controller can correct the control input often if smaller sampling times are chosen.

We have also found that in the case of modeling building thermal dynamics for small buildings with convection-based emission systems (i.e. radiator or floor heating/cooling), linear models seem to be sufficient. Moreover, the physics-based priors enforced in the regression process of the ARMAX model cannot be identified by the considered Machine Learning methods, even if abundant training data is available. These observations are potentially interesting for the commercial application of predictive control in residential buildings, where the mentioned conditions are common.

Finally, by combining data-driven methods with Robust MPC and physics-based models of some industrial components (i.e. the heat pump and water tanks) even complex control tasks, such as day-ahead reserves provision in buildings and districts can be addressed. In the case study of Chapter 5, the modeling effort of the combined system is reduced considerably, from potentially hundreds of manually-set and difficult to obtain parameters in a pure physics-based model to less than ten parameters that can be



obtained from technical sheets in our approach.

## 6.2 Future directions

The presented work takes a step towards *real-life-applicable* efficient building energy control by replacing tedious first-principles-based modeling with data-driven methods in areas where parameters are difficult to identify, i.e. building thermal dynamics. While we have demonstrated that the developed approaches work in practice, there are limitations that lead to future research directions. One immediately obvious future direction is the application of the presented models and methods to different optimization goals. Here, we have applied them to the goals of energy use reduction in Chapter 4, and frequency reserve provision in the Chapters 3 and 5. However, the presented methods are not limited to these objectives. Alternative objectives that should be tested in practice are, for example, the exploitation of time-varying energy prices, and peak shaving.

Future directions that require significantly more thinking comprise various aspects of *scalability*, which relate to increasing the boundaries of the controlled systems. For example, the numerical study in Chapter 5 has shown that more efficient optimization methods are needed if the number of controlled buildings is to be increased. Such methods should also ideally be distributed algorithms to represent the distributed architecture of the controlled energy systems. Moreover, privacy concerns related to the sharing of measurement data, as well as game-theoretic implications of distributed algorithms in energy-hub settings and related incentive structures should be considered in the optimization schemes. If, as a result of integrating a larger number of buildings, the spacial size of the considered district system is increased, also modeling the distribution network could become necessary, as thermal losses in the piping system need to be considered.

In addition to increasing the number of controlled entities, also the diversity of their types should be considered in future research: To make energy systems more efficient, sectors and energy streams will need to be coupled and combined into multi-energy-systems in the future. A control system addressing building thermal management should therefore be able to also consider a connected PV panel or to control the charging and discharging of a battery electric vehicle. It should also communicate with district heating system operators or the electricity grids to coordinate control decisions on various levels of hierarchy. While the presented work takes a step in this direction by demonstrating methods for providing electrical reserves to the grid with electrified building and district heating/cooling systems, all presented data-driven methods in this study only consider the thermal behaviour of buildings. However, in an ideal case, a data-driven modeling approach should, given a set of measurement data, independently identify different types of data (electrical loads, room temperature trajectories, states of charge, etc.) and the related entities, and then build a dynamic model of the identified multi-energy system.

This task requires a combination of classification and regression methods that has not been appropriately addressed in the literature yet. Such solutions could especially be important for legacy buildings, where labeling of data streams requires a considerable amount of manual effort when old building equipment is brought to the cloud in an effort to integrate it in advanced control schemes.

Related to this, also the field of physics-informed Machine Learning should be considered, as for different classes of data different physical priors can be applied. Chapter 4 has shown that physics-based priors can significantly improve the model accuracy and sample efficiency for models of building thermal dynamics. By embedding these priors in more flexible modeling approaches, such as the presented RF and ICNN methods, potentially also multi-energy systems could be addressed. Here, the RF and ICNN methods could potentially have an advantage compared to linear methods, as the underlying physics of multi-energy-systems contain more non-linearities compared to pure building thermal dynamics. For this task, compared to the efforts in this thesis, maintaining the balance between model flexibility and computational efficiency during operation will be even more challenging.

## A.1 Additional KPI for Case study 1 in Chapter 2




Figure A.1: Mean Squared Error (MSE) of 100 ANN and forecast correction methods.




Figure A.2: Mean Absolute Error (MAE) of 100 ANN and forecast correction methods.






Figure A.3: Coefficient of Variation of Root-Mean Squared Error (CV RMSE) of 100 ANN and forecast correction methods.


## A.2 Forecasting trajectories for Case study 2 in Chapter 2




(a)



(b)




(c)




(d)


Figure A.4: Scaled forecasting trajectories on the training set for buildings (a)-(d).




(a)



(b)



(c)



(d)

Figure A.5: Scaled forecasting trajectories on the test set for buildings (a)-(d).

## A.3 Additional Figures for Chapter 4




Figure A.6: Example from the validation set of the one-hot solar model. The blue line denotes the true solar gains through a window. The orange line denotes the predicted gains by the model.




Figure A.7: Baseline controller example for the heating case in bedroom 2. (a): Temperature in the bedroom. The limits of the hysteresis controller are shown in dashed black. (b): Relative control input, i.e. the fraction of time where the maximum control input is applied during one control step. (c): Measured ambient temperature at the experiment site. (d): Global solar irradiation at the experiment site.




Figure A.8: Temperature analysis of MPC vs. baseline controller for heating experiments in bedroom 1.



# Bibliography

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X., Brain, G., 2016. TensorFlow: A system for large-scale machine learning, in: Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16), pp. 265–283.
- Abergel, T., Dulac, J., Hamilton, I., Jordan, M., Pradeep, A., 2019. Global status report for buildings and construction—towards a zero-emissions, efficient and resilient buildings and construction sector.
- Afram, A., Janabi-Sharifi, F., 2014. Theory and applications of HVAC control systems - A review of model predictive control (MPC). *Building and Environment* 72, 343–355.
- Aggarwal, C.C., 2018. Text sequence modeling and deep learning, in: *Machine Learning for Text*. Springer International Publishing, Cham, pp. 305–360.
- Ahmad, T., Chen, H., Guo, Y., Wang, J., 2018. A comprehensive overview on the data driven and large scale based approaches for forecasting of building energy demand: A review. *Energy and Buildings* 165, 301–320.
- Alahäivälä, A., Corbishley, J., Ekström, J., Jokisalo, J., Lehtonen, M., 2017. A control framework for the utilization of heating load flexibility in a day-ahead market. *Electric Power Systems Research* 145, 44–54.
- Amasyali, K., El-Gohary, N.M., 2018. A review of data-driven building energy consumption prediction studies. *Renewable and Sustainable Energy Reviews* 81, 1192–1205.
- Amos, B., Xu, L., Kolter, J.Z., 2017. Input Convex Neural Networks, in: 34th International Conference on Machine Learning, ICML 2017, PMLR. pp. 192–206.
- Andersen, M.S., Dahl, J., Vandenberghe, L., 2004. CVXOPT. URL: <http://cvxopt.org/>. Accessed 13 Jan 2021.
- Arnold, M., Andersson, G., 2011. Model predictive control of energy storage including uncertain forecasts. *Proceedings of the 17th Power Systems Computation Conference (PSCC)* 22, 60–71.

- Arnold, M., Negenborn, R.R., Andersson, G., De Schutter, B., 2009. Model-based predictive control applied to multi-carrier energy systems, in: 2009 IEEE Power and Energy Society General Meeting, PES '09, pp. 1–8.
- Aswani, A., Master, N., Taneja, J., Krioukov, A., Culler, D., Tomlin, C., 2012. Energy-efficient building HVAC control using hybrid system LBMPC. IFAC Proceedings Volumes (IFAC-PapersOnline) 45, 496–501.
- Basheer, I., Hajmeer, M., 2000. Artificial neural networks: fundamentals, computing, design, and application. *Journal of Microbiological Methods* 43, 3–31.
- Belimo, 2021. Belimo Energy Valve™ Technical Documentation. URL: [https://www.belimo.com/mam/americas/technical\\_documents/Techdocs/energy\\_valve\\_technical\\_document.pdf](https://www.belimo.com/mam/americas/technical_documents/Techdocs/energy_valve_technical_document.pdf). Accessed 27 May 2021.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., Nemirovski, A., 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* 99, 351–376.
- Bengio, Y., 2012. Practical recommendations for gradient-based training of deep architectures, in: Montavon, G., Orr, G., Müller, K.R. (Eds.), *Neural Networks: Tricks of the Trade*. 2 ed.. Springer Berlin Heidelberg, pp. 437—478.
- Bescond, P.L., 2020. Cyclical features encoding, it's about time! - towards data science. URL: <https://towardsdatascience.com/cyclical-features-encoding-its-about-time-ce23581845ca>. Accessed 26 May 2021.
- Blum, D.H., Xu, N., Norford, L.K., 2016. A novel multi-market optimization problem for commercial heating, ventilation, and air-conditioning systems providing ancillary services using multi-zone inverse comprehensive room transfer functions. *Science and Technology for the Built Environment* 22, 783–797.
- Borrelli, F., Bemporad, A., Morari, M., 2017. *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press.
- Borsche, T., Oldewurtel, F., Andersson, G., 2014. Scenario-based MPC for energy schedule compliance with demand response. IFAC Proceedings Volumes (IFAC-PapersOnline) 19, 10299–10304.
- Brownlee, J., 2018. Why one-hot encode data in machine learning? URL: <https://machinelearningmastery.com/why-one-hot-encode-data-in-machine-learning/>. Accessed 25 Oct 2018.
- Brück, D., Elmqvist, H., Erik Mattsson, S., Olsson, H., 2002. Dymola for multi-engineering modeling and simulation, in: 2nd International Modelica Conference, pp. 1–8.
- Bünning, F., Heer, P., Smith, R.S., Lygeros, J., 2020a. Improved day ahead heating demand forecasting by online correction methods. *Energy and Buildings* 211, 109821.

- Bünning, F., Huber, B., Heer, P., AbouDonia, A., Lygeros, J., 2020b. Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings. *Energy and Buildings* 211, 109792.
- Bünning, F., Huber, B., Schalbetter, A., AbouDonia, A., Hudoba de Badyn, M., Heer, P., Smith, R.S., Lygeros, J., 2021a. Physics-informed linear regression is competitive with two Machine Learning methods in residential building MPC. Accepted at *Applied Energy*.
- Bünning, F., Schalbetter, A., AbouDonia, A., Hudoba de Badyn, M., Heer, P., Lygeros, J., 2021b. Input convex neural networks for building MPC, in: *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, PMLR. pp. 251–262.
- Bünning, F., Warrington, J., Heer, P., Smith, R.S., Lygeros, J., 2020c. Frequency regulation with heat pumps using robust MPC with affine policies. *IFAC-PapersOnLine* 53, 13210–13215.
- Bünning, F., Warrington, J., Heer, P., Smith, R.S., Lygeros, J., 2020d. Robust MPC with data-driven demand forecasting for frequency regulation with heat pumps. [arXiv:2009.06920](https://arxiv.org/abs/2009.06920).
- Bünning, F., Wetter, M., Fuchs, M., Müller, D., 2018. Bidirectional low temperature district energy systems with agent-based control: Performance comparison and operation optimization. *Applied Energy* 209, 502–515.
- Castilla, M., Álvarez, J.D., Normey-Rico, J.E., Rodríguez, F., 2014. Thermal comfort control using a non-linear MPC strategy: A real case of study in a bioclimatic building. *Journal of Process Control* 24, 703–713.
- Chen, C., Wang, J., Heo, Y., Kishore, S., 2013. MPC-based appliance scheduling for residential building energy management controller. *IEEE Transactions on Smart Grid* 4, 1401–1410.
- Chen, X., Sim, M., Sun, P., Zhang, J., 2008. A linear decision-based approximation approach to stochastic programming. *Operations Research* 56, 344–357.
- Chen, Y., Shi, Y., Zhang, B., 2019. Optimal control via neural networks: A convex approach, in: *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–21.
- Cholewa, T., Siuta-Olcha, A., Smolarz, A., Muryjas, P., Wolszczak, P., Guz, Ł., Balaras, C.A., 2021. On the short term forecasting of heat power for heating of building. *Journal of Cleaner Production* 307, 127232.
- Chollet, F., 2018. Keras: The Python Deep Learning library. *Astrophysics Source Code Library* URL: <http://adsabs.harvard.edu/abs/2018ascl.soft06022C>. Accessed 25 Oct 2018.
- Coulson, J., Lygeros, J., Dorfler, F., 2019. Data-enabled predictive control: In the shallows of the DeePC, in: *2019 18th European Control Conference, ECC 2019*, pp. 307–312.

- Darivianakis, G., Georghiou, A., Smith, R.S., Lygeros, J., 2017. The power of diversity: Data-driven robust predictive control for energy-efficient buildings and districts. *IEEE Transactions on Control Systems Technology* 99, 132–145.
- De Coninck, R., Magnusson, F., Åkesson, J., Helsen, L., 2015. Toolbox for development and validation of grey-box building models for forecasting and control. *Journal of Building Performance Simulation* 9, 288–303.
- De Felice, M., Yao, X., 2011. Short-term load forecasting with neural network ensembles: A comparative study. *IEEE Computational Intelligence Magazine* 6, 47–56.
- Dimitriou, V., Firth, S.K., Hassan, T.M., Kane, T., Coleman, M., 2015. Data-driven simple thermal models: The importance of the parameter estimates. *Energy Procedia* 78, 2614–2619.
- Drgoňa, J., Arroyo, J., Cupeiro Figueroa, I., Blum, D., Arendt, K., Kim, D., Ollé, E.P., Oravec, J., Wetter, M., Vrabie, D.L., Helsen, L., 2020a. All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50, 190–232.
- Drgoňa, J., Tuor, A.R., Chandan, V., Vrabie, D.L., 2020b. Physics-constrained deep learning of multi-zone building thermal dynamics. *arXiv:2011.05987*.
- Efron, B., 1979. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* 7, 1 – 26.
- Empa, 2018. Communication - EHUB. URL: <https://www.empa.ch/web/s604/ehub>. Accessed 18 Mar 2021.
- Ferkl, L., Jan Široký, 2010. Ceiling radiant cooling: Comparison of ARMAX and subspace identification modelling methods. *Building and Environment* 45, 205–212.
- Ferreira, P., Ruano, A., Silva, S., Conceição, E., 2012. Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings* 55, 238–251.
- Finck, C., Li, R., Zeiler, W., 2019. Economic model predictive control for demand flexibility of a residential building. *Energy* 176, 365–379.
- Fischer, D., Madani, H., 2017. On heat pumps in smart grids: A review. *Renewable and Sustainable Energy Reviews* 70, 342–357.
- Flamm, B., Eichler, A., Warrington, J., Lygeros, J., 2021. Two-stage dual dynamic programming with application to nonlinear hydro scheduling. *IEEE Transactions on Control Systems Technology* 29, 96–107.
- Foucquier, A., Robert, S., Suard, F., Stéphan, L., Jay, A., 2013. State of the art in building modelling and energy performances prediction: A review. *Renewable and Sustainable Energy Reviews* 23, 272–288.

- Geidl, M., Arnoux, B., Plaisted, T., Dufour, S., 2017. A fully operational virtual energy storage network providing flexibility for the power system, in: 12th IEA Heat Pump Conference, pp. 15–18.
- Geidl, M., Koeppel, G., Favre-Perrod, P., 2007. The Energy Hub—A powerful concept for future energy systems, in: Third annual Carnegie mellon conference on the electricity industry, pp. 1–10.
- Gelazanskas, L., Gamage, K.A., 2014. Demand side management in smart grid: A review and proposals for future direction. *Sustainable Cities and Society* 11, 22–30.
- Georghiou, A., Wiesemann, W., Kuhn, D., 2015. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming* 152, 301–338.
- Gholamzadehmir, M., Del Pero, C., Buffa, S., Fedrizzi, R., Aste, N., 2020. Adaptive-predictive control strategy for HVAC systems in smart buildings – A review. *Sustainable Cities and Society* 63, 102480.
- Goh, J., Sim, M., 2011. Robust optimization made easy with ROME. *Operations Research* 59, 973–985.
- Goulart, P.J., Kerrigan, E.C., Maciejowski, J.M., 2006. Optimization over state feedback policies for robust control with constraints. *Automatica* 42, 523–533.
- Graser, K., Adel, A., Baur, M., Pont, D.S., Thoma, A., 2021. Parallel paths of inquiry: Detailing for DFAB HOUSE. *Technology Architecture and Design* 5, 38–43.
- Hameed Shaikh, P., Bin Mohd Nor, N., Nallagownden, P., Elamvazuthi, I., Ibrahim, T., 2014. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews* 34, 409–429.
- Hammer, B., Gersmann, K., 2003. A note on the universal approximation capability of support vector machines. *Neural Processing Letters* 17, 43–53.
- Hansen, N., Müller, S.D., Koumoutsakos, P., 2003. Reducing the time complexity of the de-randomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation* 11, 1–18.
- Hao, H., Kowli, A., Lin, Y., Barooah, P., Meyn, S., 2013. Ancillary service for the grid via control of commercial building HVAC systems, in: Proceedings of the American Control Conference, pp. 467–472.
- Hao, H., Sanandaji, B.M., Poolla, K., Vincent, T.L., 2015. Potentials and economics of residential thermal loads providing regulation reserve. *Energy Policy* 79, 115–126.
- Harild Rasmussen, T.B., Wu, Q., Zhang, M., 2021. Primary frequency support from local control of large-scale heat pumps. *International Journal of Electrical Power & Energy Systems* 133, 107270.

- Harish, V.S., Kumar, A., 2016. A review on modeling and simulation of building energy systems. *Renewable and Sustainable Energy Reviews* 56, 1272–1292.
- van der Heijde, B., Vandermeulen, A., Salenbien, R., Helsen, L., 2019. Integrated optimal design and control of fourth generation district heating networks with thermal energy storage. *Energies* 12, 2766.
- Heisel, F., Hebel, D.E., Sobek, W., 2019. Resource-respectful construction – the case of the Urban Mining and Recycling unit (UMAR). *IOP Conference Series: Earth and Environmental Science* 225, 012049.
- Heisel, F., Rau-Oberhuber, S., 2020. Calculation and evaluation of circularity indicators for the built environment using the case studies of UMAR and Madaster. *Journal of Cleaner Production* 243, 118482.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D., 2018. Deep reinforcement learning that matters, in: 32nd AAAI Conference on Artificial Intelligence, AAAI 2018, pp. 3207–3214.
- Hennessy, J., Li, H., Wallin, F., Thorin, E., 2019. Flexibility in thermal grids: A review of short-term storage in district heating distribution networks. *Energy Procedia* 158, 2430–2434.
- Henze, G.P., 2013. Model predictive control for buildings: a quantum leap? *Journal of Building Performance Simulation* 6, 157–158.
- Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., Slaney, M., Weiss, R.J., Wilson, K., 2017. CNN architectures for large-scale audio classification, in: ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 131–135.
- Hilliard, T., Swan, L., Qin, Z., 2017. Experimental implementation of whole building MPC with zone based thermal comfort adjustments. *Building and Environment* 125, 326–338.
- Hohmann, M., Warrington, J., Lygeros, J., 2019. A two-stage polynomial approach to stochastic optimization of district heating networks. *Sustainable Energy, Grids and Networks* 17, 100177.
- Hornik, K., 1991. Approximation capabilities of multilayer feedforward networks. *Neural Networks* 4, 251–257.
- Huang, J., Smola, A.J., Gretton, A., Borgwardt, K.M., Schölkopf, B., 2006. Correcting sample selection bias by unlabeled data, in: NIPS’06 Proceedings of the 19th International Conference on Neural Information Processing Systems, pp. 601–608.
- Ilić, D.D., 2020. Classification of measures for dealing with district heating load variations—a systematic review. *Energies* 14, 3.

- Ivanova, P., Sauhats, A., Linkevics, O., 2019. District heating technologies: Is it chance for CHP plants in variable and competitive operation conditions? *IEEE Transactions on Industry Applications* 55, 35–42.
- Jain, A., Smarra, F., Reticcioli, E., D’Innocenzo, A., Morari, M., 2020. NeurOpt: Neural network based optimization for building energy management and climate control, in: *Proceedings of Machine Learning Research*, PMLR. pp. 445–454.
- Jamieson, K., Talwalkar, A., 2016. Non-stochastic best arm identification and hyperparameter optimization, in: *19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 240–248.
- Jetcheva, J.G., Majidpour, M., Chen, W.P., 2014. Neural network model ensembles for building-level electricity load forecasts. *Energy and Buildings* 84, 214–223.
- Johansson, C., Bergkvist, M., Geysen, D., Somer, O.D., Lavesson, N., Vanhoudt, D., 2017. Operational demand forecasting in district heating systems using ensembles of online machine learning algorithms. *Energy Procedia* 116, 208–216.
- Johnson, S.C., Papageorgiou, D.J., Mallapragada, D.S., Deetjen, T.A., Rhodes, J.D., Webber, M.E., 2019. Evaluating rotational inertia as a component of grid reliability with high penetrations of variable renewable energy. *Energy* 180, 258–271.
- Jovanović, R., Sretenović, A.A., Živković, B.D., 2015. Ensemble of various neural networks for prediction of heating energy consumption. *Energy and Buildings* 94, 189–199.
- Junker, R.G., Azar, A.G., Lopes, R.A., Lindberg, K.B., Reynders, G., Relan, R., Madsen, H., 2018. Characterizing the energy flexibility of buildings and districts. *Applied Energy* 225, 175–182.
- Kakkos, E., Heisel, F., Hebel, D.E., Hischer, R., 2019. Environmental assessment of the Urban Mining and Recycling (UMAR) unit by applying the LCA framework. *IOP Conference Series: Earth and Environmental Science* 225, 012043.
- Kathirgamanathan, A., De Rosa, M., Mangina, E., Finn, D.P., 2021. Data-driven predictive control for unlocking building energy flexibility: A review. *Renewable and Sustainable Energy Reviews* 135, 110120.
- Kathirgamanathan, A., Péan, T., Zhang, K., De Rosa, M., Salom, J., Kummert, M., Finn, D.P., 2020. Towards standardising market-independent indicators for quantifying energy flexibility in buildings. *Energy and Buildings* 220, 110027.
- Kato, K., Sakawa, M., Keiichi Ishimaru, Ushiro, S., Toshihiro Shibano, 2008. Heat load prediction through recurrent neural network in district heating and cooling systems, in: *2008 IEEE International Conference on Systems, Man and Cybernetics*, pp. 1401–1406.
- Kim, Y.J., Blum, D.H., Xu, N., Su, L., Norford, L.K., 2016. Technologies and magnitude of ancillary services provided by commercial buildings. *Proceedings of the IEEE* 104, 758–779.

- Kim, Y.J., Norford, L.K., Kirtley, J.L., 2015. Modeling and analysis of a variable speed heat pump for frequency regulation through direct load control. *IEEE Transactions on Power Systems* 30, 397–408.
- Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- Koch, S., Mathieu, J.L., Callaway, D.S., 2011. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services, in: 17th Power Systems Computation Conference, PSCC 2011, pp. 1–7.
- Kondoh, J., Lu, N., Hammerstrom, D.J., 2011. An evaluation of the water heater load potential for providing regulation service. *IEEE Transactions on Power Systems* 26, 1309–1316.
- Kouvaritakis, B., Cannon, M., 2016. Model Predictive Control - Classical, Robust and Stochastic. Springer International Publishing.
- Kwok, S.S., Lee, E.W., 2011. A study of the importance of occupancy to building cooling load in prediction by intelligent approach. *Energy Conversion and Management* 52, 2555–2564.
- Langevin, J., Harris, C.B., Satre-Meloy, A., Chandra-Putra, H., Speake, A., Present, E., Adhikari, R., Wilson, E.J., Satchwell, A.J., 2021. US building energy efficiency and flexibility as an electric grid resource. *Joule* 5, 2102–2128.
- Lee, S., Karava, P., 2020. Towards smart buildings with self-tuned indoor thermal environments – A critical review. *Energy and Buildings* 224, 110172.
- Lefebure, N., Khosravi, M., Hudoba de Badyn, M., Bünning, F., Lygeros, J., Jones, C., Smith, R.S., 2021. Distributed model predictive control of buildings and energy hubs. [arXiv:1907.04490](https://arxiv.org/abs/1907.04490).
- Leitner, S.H., Mahnke, W., 2006. OPC UA-service-oriented architecture for industrial applications. *ABB Corporate Research Center* 48, 22.
- Leung, M., Tse, N.C., Lai, L., Chow, T., 2012. The use of occupancy space electrical power demand in building cooling load prediction. *Energy and Buildings* 55, 151–163.
- Li, Y., O’Neill, Z., Zhang, L., Chen, J., Im, P., DeGraw, J., 2021. Grey-box modeling and application for building energy simulations - A critical review. *Renewable and Sustainable Energy Reviews* 146, 111174.
- Li, Z., Wu, W., Shahidehpour, M., Wang, J., Zhang, B., 2016. Combined heat and power dispatch considering pipeline energy storage of district heating network. *IEEE Transactions on Sustainable Energy* 7, 12–22.
- Lin, Y., Barooah, P., Meyn, S.P., 2013. Low-frequency power-grid ancillary services from commercial building HVAC systems, in: 2013 IEEE International Conference on Smart Grid Communications, pp. 169–174.



- Löfberg, J., 2004. YALMIP: A toolbox for modeling and optimization in MATLAB, in: Proceedings of the IEEE International Symposium on Computer-Aided Control System Design, pp. 284–289.
- Löfberg, J., 2012. Automatic robust convex programming. *Optimization Methods and Software* 27, 115–129.
- Loshchilov, I., Hutter, F., 2016. CMA-ES for hyperparameter optimization of deep neural networks.
- Louppe, G., 2014. Understanding random forests: From theory to practice. [arXiv:1407.7502](https://arxiv.org/abs/1407.7502).
- Lund, H., Duic, N., Østergaard, P.A., Mathiesen, B.V., 2018. Future district heating systems and technologies: On the role of smart energy systems and 4th generation district heating. *Energy* 165, 614–619.
- Lund, H., Østergaard, P.A., Nielsen, T.B., Werner, S., Thorsen, J.E., Gudmundsson, O., Arabkoohsar, A., Mathiesen, B.V., 2021. Perspectives on fourth and fifth generation district heating. *Energy* 227, 120520.
- Lutter, M., Ritter, C., Peters, J., 2019. Deep lagrangian networks: Using physics as model prior for deep learning. [arXiv:1907.04490](https://arxiv.org/abs/1907.04490).
- Lydon, G.P., Hofer, J., Svetozarevic, B., Nagy, Z., Schlueter, A., 2017. Coupling energy systems with lightweight structures for a net plus energy building. *Applied Energy* 189, 310–326.
- Lygnerud, K., Ottosson, J., Kensby, J., Johansson, L., 2021. Business models combining heat pumps and district heating in buildings generate cost and emission savings. *Energy* 234, 121202.
- Ma, J., Qin, S.J., Salsbury, T., 2014. Application of economic MPC to the energy and demand minimization of a commercial building. *Journal of Process Control* 24, 1282–1291.
- Maasoumy, M., Ortiz, J., Culler, D., Sangiovanni-Vincentelli, A., 2013. Flexibility of commercial building HVAC fan as ancillary service for smart grid, in: Proceedings of Green Energy and Systems Conference 2013.
- Maddalena, E.T., Lian, Y., Jones, C.N., 2020. Data-driven methods for building control — A review and promising future directions. *Control Engineering Practice* 95, 104211.
- Mai, W., Chung, C.Y., 2015. Economic MPC of aggregating commercial buildings for providing flexible power reserve. *IEEE Transactions on Power Systems* 30, 2685–2694.
- Manek, G., Kolter, J.Z., 2020. Learning Stable Deep Dynamics Models. [arXiv:2001.06116](https://arxiv.org/abs/2001.06116).
- Mania, H., Guy, A., Recht, B., 2018. Simple random search of static linear policies is competitive for reinforcement learning, in: 32nd Conference on Neural Information Processing Systems (NIPS 2018), pp. 1805–1814.

- Mariano-Hernández, D., Hernández-Callejo, L., Zorita-Lamadrid, A., Duque-Pérez, O., Santos García, F., 2021. A review of strategies for building energy management system: Model predictive control, demand side management, optimization, and fault detect & diagnosis. *Journal of Building Engineering* 33, 101692.
- Markovsky, I., Willems, J.C., Van Huffel, S., De Moor, B., 2006. Exact and approximate modeling of linear systems: A behavioral approach. Society for Industrial and Applied Mathematics.
- Márquez-Neila, P., Salzmann, M., Fua, P., 2017. Imposing hard constraints on deep networks: Promises and limitations. [arXiv:1706.02025](https://arxiv.org/abs/1706.02025).
- Mat Daut, M.A., Hassan, M.Y., Abdullah, H., Rahman, H.A., Abdullah, M.P., Hussin, F., 2017. Building electrical energy consumption forecasting analysis using conventional and artificial intelligence methods: A review. *Renewable and Sustainable Energy Reviews* 70, 1108–1118.
- Mattsson, S.E., Elmqvist, H., Otter, M., 1998. Physical system modeling with Modelica. *Control Engineering Practice* 6, 501–510.
- Mckinney, W., 2011. pandas: a foundational python library for data analysis and statistics, in: *Python for High Performance and Scientific Computing*, pp. 1–9.
- Mestekemper, T., Kauermann, G., Smith, M.S., 2013. A comparison of periodic autoregressive and dynamic factor models in intraday energy demand forecasting. *International Journal of Forecasting* 29, 1–12.
- Mocanu, E., Mocanu, D.C., Nguyen, P.H., Liotta, A., Webber, M.E., Gibescu, M., Slootweg, J.G., 2019. On-line building energy optimization using deep reinforcement learning. *IEEE Transactions on Smart Grid* 10, 3698–3708.
- Morari, M., H. Lee, J., 1999. Model predictive control: Past, present and future. *Computers and Chemical Engineering* 23, 667–682.
- Mugnini, A., Coccia, G., Polonara, F., Arteconi, A., 2020. Performance assessment of data-driven and physical-based models to predict building energy demand in model predictive controls. *Energies* 13, 3125.
- Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted Boltzmann machines, in: *Proceedings of the 27th international conference on machine learning*, pp. 807–814.
- Olama, M., Kuruganti, T., Nutaro, J., Dong, J., 2018. Coordination and control of building HVAC systems to provide frequency regulation to the electric grid. *Energies* 11, 1852.
- Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Wirth, K., 2010. Energy efficient building climate control using stochastic model predictive control and weather predictions, in: *Proceedings of the 2010 American Control Conference*, pp. 5100–5105.

- Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Morari, M., 2012. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* 45, 15–27.
- Oldewurtel, F., Sturzenegger, D., Andersson, G., Morari, M., Smith, R.S., 2013. Towards a standardized building assessment for demand response, in: *Proceedings of the IEEE Conference on Decision and Control*, pp. 7083–7088.
- Pagh Nielsen, M., Sørensen, K., 2021. Dynamic modeling of heat pumps for ancillary services in local district heating concepts, in: *Proceedings of The 61st SIMS Conference on Simulation and Modelling SIMS 2020*, pp. 39–46.
- Palensky, P., Dietrich, D., 2011. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics* 7, 381–388.
- Park, T.C., Kim, U.S., Kim, L.H., Jo, B.W., Yeo, Y.K., 2010. Heat consumption forecasting using partial least squares, artificial neural network and support vector regression techniques in district heating systems. *Korean Journal of Chemical Engineering* 27, 1063–1071.
- Paudel, S., Elmtiri, M., Kling, W.L., Corre, O.L., Lacarrière, B., 2014. Pseudo dynamic transitional modeling of building heating energy demand using artificial neural network. *Energy and Buildings* 70, 81–93.
- Péan, T.Q., Salom, J., Costa-Castelló, R., 2019. Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings. *Journal of Process Control* 74, 35–49.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, É., 2011. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research* 12, 2825–2830.
- Picard, D., Jorissen, F., Helsen, L., 2015. Methodology for obtaining linear state space building energy simulation models, in: *Proceedings of the 11th International Modelica Conference*, pp. 51–58.
- Picard, D., Sourbron, M., Jorissen, F., Cigler, J., Vána, Z., 2016. Comparison of model predictive control performance using grey-box and white-box controller models of a multi-zone office building, in: *4th International High Performance Buildings Conference*.
- Pinto, G., Piscitelli, M.S., Vázquez-Canteli, J.R., Nagy, Z., Capozzoli, A., 2021. Coordinated energy management for a cluster of buildings through deep reinforcement learning. *Energy* 229, 120725.
- PJM, 2019. PJM Manual 12: Balancing Operations. URL: <https://www.pjm.com/-/media/documents/manuals/m12-redline.ashx?la=en>. Accessed 23 Jun 2019.

- Powell, M.J.D., 1994. A direct search optimization method that models the objective and constraint functions by linear interpolation, in: *Advances in Optimization and Numerical Analysis*. Springer Netherlands, pp. 51–67.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2017. Physics informed deep learning (Part II): Data-driven discovery of nonlinear partial differential equations. *arXiv:1711.10566*.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707.
- Rastegarpour, S., Ferrarini, L., Gros, S., 2020a. Economic NMPC for multiple buildings connected to a heat pump and thermal and electrical storages. *IFAC-PapersOnLine* 53, 17089–17094.
- Rastegarpour, S., Gros, S., Ferrarini, L., 2020b. MPC approaches for modulating air-to-water heat pumps in radiant-floor buildings. *Control Engineering Practice* 95, 104209.
- Recht, B., 2018. Updates on Policy Gradients – arg min blog. URL: <http://www.argmin.net/2018/03/13/pg-saga/>. Accessed 23 Oct 2018.
- Richner, P., Heer, P., Largo, R., Marchesi, E., Zimmermann, M., 2017. NEST - A platform for the acceleration of innovation in buildings. *Informes de la Construcción* 69, e222.
- Rockett, P., Hathway, E.A., 2017. Model-predictive control for non-domestic buildings: a critical review and prospects. *Building Research & Information* 45, 556–571.
- Romero Rodríguez, L., Brennenstuhl, M., Yadack, M., Boch, P., Eicker, U., 2019. Heuristic optimization of clusters of heat pumps: A simulation and case study of residential frequency reserve. *Applied Energy* 233-234, 943–958.
- Rominger, J., Kern, F., Schmeck, H., 2018. Provision of frequency containment reserve with an aggregate of air handling units, in: *Computer Science - Research and Development*, Springer Verlag. pp. 215–221.
- Saloux, E., Candanedo, J.A., 2018. Forecasting district heating demand using machine learning algorithms. *Energy Procedia* 149, 59–68.
- Salpakari, J., Mikkola, J., Lund, P.D., 2016. Improved flexibility with large-scale variable renewable power in cities through optimal demand side management and power-to-heat conversion. *Energy Conversion and Management* 126, 649–661.
- Schalbetter, A., 2020. Input Convex Neural Networks for Energy Optimization in an occupied Apartment. ETH Zurich master thesis.
- ScienceDirect, 2021. Heating Degree Day - an overview | ScienceDirect Topics. URL: <https://www.sciencedirect.com/topics/engineering/heating-degree-day>. Accessed 08 Mar 2021.

- Serale, G., Fiorentini, M., Capozzoli, A., Bernardini, D., Bemporad, A., 2018. Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities. *Energies* 11, 631.
- Shanmuganathan, S., 2016. Artificial neural network modelling: An introduction, in: *Artificial neural network modelling*. Springer, pp. 1–14.
- Sharma, R., Farimani, A.B., Gomes, J., Eastman, P., Pande, V., 2018. Weakly-supervised deep learning of heat transport via physics informed loss. [arXiv:1807.11374](https://arxiv.org/abs/1807.11374).
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Van Den Driessche, G., Graepel, T., Hassabis, D., 2017. Mastering the game of Go without human knowledge. *Nature* 550, 354–359.
- Široký, J., Oldewurtel, F., Cigler, J., Prívvara, S., 2011. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy* 88, 3079–3087.
- Smarra, F., Jain, A., de Rubeis, T., Ambrosini, D., D’Innocenzo, A., Mangharam, R., 2018. Data-driven model predictive control using random forests for building energy optimization and climate control. *Applied Energy* 226, 1252–1272.
- Sturzenegger, D., Gyalistras, D., Morari, M., Smith, R.S., 2016. Model predictive climate control of a swiss office building: Implementation, results, and cost–benefit analysis. *IEEE Transactions on Control Systems Technology* 24, 1–12.
- Sturzenegger, D., Gyalistras, D., Semeraro, V., Morari, M., Smith, R.S., 2014. BRCM Matlab toolbox: Model generation for model predictive building control, in: *Proceedings of the American Control Conference*, pp. 1063–1069.
- Suganthi, L., Samuel, A.A., 2012. Energy models for demand forecasting—A review. *Renewable and Sustainable Energy Reviews* 16, 1223–1240.
- Sundström, C., Jung, D., Blom, A., 2017. Analysis of optimal energy management in smart homes using MPC, in: *2016 European Control Conference, ECC 2016*, pp. 2066–2071.
- Suryanarayana, G., Lago, J., Geysen, D., Aleksiejuk, P., Johansson, C., 2018. Thermal load forecasting in district heating networks using deep learning and advanced feature selection methods. *Energy* 157, 141–149.
- Tang, H., Wang, S., 2021. Energy flexibility quantification of grid-responsive buildings: Energy flexibility index and assessment of their effectiveness for applications. *Energy* 221, 119756.
- Terreros, O., Spreitzhofer, J., Basciotti, D., Schmidt, R.R., Esterl, T., Pober, M., Kerschbaumer, M., Ziegler, M., 2020. Electricity market options for heat pumps in rural district heating networks in Austria. *Energy* 196, 116875.

- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 267–288.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P., 2017. Domain randomization for transferring deep neural networks from simulation to the real world, in: *IEEE International Conference on Intelligent Robots and Systems*, pp. 23–30.
- Touretzky, C.R., Baldea, M., 2014. Integrating scheduling and control for economic MPC of buildings with energy storage. *Journal of Process Control* 24, 1292–1300.
- Vandermeulen, A., van der Heijde, B., Helsen, L., 2018. Controlling district heating and cooling networks to unlock flexibility: A review. *Energy* 151, 103–115.
- Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, I., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., Vijaykumar, A., Bardelli, A.P., Rothberg, A., Hilboll, A., Kloeckner, A., Scopatz, A., Lee, A., Rokem, A., Woods, C.N., Fulton, C., Masson, C., Häggström, C., Fitzgerald, C., Nicholson, D.A., Hagen, D.R., Pasechnik, D.V., Olivetti, E., Martin, E., Wieser, E., Silva, F., Lenders, F., Wilhelm, F., Young, G., Price, G.A., Ingold, G.L., Allen, G.E., Lee, G.R., Audren, H., Probst, I., Dietrich, J.P., Silterra, J., Webber, J.T., Slavič, J., Nothman, J., Buchner, J., Kulick, J., Schönberger, J.L., de Miranda Cardoso, J.V., Reimer, J., Harrington, J., Rodríguez, J.L.C., Nunez-Iglesias, J., Kuczynski, J., Tritz, K., Thoma, M., Newville, M., Kümmerer, M., Bolingbroke, M., Tartre, M., Pak, M., Smith, N.J., Nowaczyk, N., Shebanov, N., Pavlyk, O., Brodtkorb, P.A., Lee, P., McGibbon, R.T., Feldbauer, R., Lewis, S., Tygier, S., Sievert, S., Vigna, S., Peterson, S., More, S., Pudlik, T., Oshima, T., Pingel, T.J., Robitaille, T.P., Spura, T., Jones, T.R., Cera, T., Leslie, T., Zito, T., Krauss, T., Upadhyay, U., Halchenko, Y.O., Vázquez-Baeza, Y., 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* 17, 261–272.
- Vrettos, E., Kara, E.C., MacDonald, J., Andersson, G., Callaway, D.S., 2018a. Experimental demonstration of frequency regulation by commercial buildings-part I: Modeling and hierarchical control design. *IEEE Transactions on Smart Grid* 9, 3213–3223.
- Vrettos, E., Kara, E.C., MacDonald, J., Andersson, G., Callaway, D.S., 2018b. Experimental demonstration of frequency regulation by commercial buildings-part II: Results and performance evaluation. *IEEE Transactions on Smart Grid* 9, 3224–3234.
- Vrettos, E., Oldewurtel, F., Andersson, G., 2016. Robust Energy-Constrained Frequency Reserves from Aggregations of Commercial Buildings. *IEEE Transactions on Power Systems* 31, 4272–4285.
- Wang, H., Wang, S., Shan, K., 2020. Experimental study on the dynamics, quality and impacts of using variable-speed pumps in buildings for frequency regulation of smart power grids. *Energy* 199, 117406.

- Wang, J., Li, S., Chen, H., Yuan, Y., Huang, Y., 2019. Data-driven model predictive control for building climate control: Three case studies on different buildings. *Building and Environment* 160, 106204.
- Wang, S., Xu, X., 2006. Simplified building model for transient thermal performance estimation using GA-based parameter identification. *International Journal of Thermal Sciences* 45, 419–432.
- Wang, Z., Hong, T., 2020. Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy* 269, 115036.
- Wang, Z., Srinivasan, R.S., 2017. A review of artificial intelligence based building energy use prediction: Contrasting the capabilities of single and ensemble prediction models. *Renewable and Sustainable Energy Reviews* 75, 796–808.
- Warrington, J., Goulart, P.J., Mariethoz, S., Morari, M., 2012. Robust reserve operation in power systems using affine policies, in: 51st IEEE Conference on Decision and Control (CDC), pp. 1111–1117.
- Warrington, J., Hohl, C., Goulart, P.J., Morari, M., 2014. Optimal unit commitment accounting for robust affine reserve policies, in: 2014 American Control Conference, pp. 5049–5055.
- Wirtz, M., Kivilip, L., Remmen, P., Müller, D., 2020. 5th generation district heating: A novel design approach based on mathematical optimization. *Applied Energy* 260, 114158.
- Wojna, Z., Gorban, A.N., Lee, D.S., Murphy, K., Yu, Q., Li, Y., Ibarz, J., 2018. Attention-based extraction of structured information from street view imagery, in: Proceedings of the International Conference on Document Analysis and Recognition, ICDAR, pp. 844–850.
- Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, Ł., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J., 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. [arXiv:1609.08144](https://arxiv.org/abs/1609.08144).
- Xu, X., Lyu, Q., Qadrdan, M., Wu, J., 2020. Quantification of flexibility of a district heating system for the power grid. *IEEE Transactions on Sustainable Energy* 11, 2617–2630.
- Yang, S., Wan, M.P., Chen, W., Ng, B.F., Dubey, S., 2020. Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. *Applied Energy* 271, 115147.
- Yudong Ma, Borrelli, F., Hency, B., Coffey, B., Benghea, S., Haves, P., 2012. Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology* 20, 796–803.

- Žáčková, E., Váňa, Z., Cigler, J., 2014. Towards the real-life implementation of MPC for an office building: Identification issues. *Applied Energy* 135, 53–62.
- Zhang, X., Kamgarpour, M., Georghiou, A., Goulart, P., Lygeros, J., 2017. Robust optimal control with adjustable uncertainty sets. *Automatica* 75, 249–259.
- Zhang, Y., Zhang, T., Wang, R., Liu, Y., Guo, B., 2015. Optimal operation of a smart residential microgrid based on model predictive control by considering uncertainties and storage impacts. *Solar Energy* 122, 1052–1065.
- Zhao, H.X., Magoulès, F., 2012. A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews* 16, 3586–3592.
- Zhao, P., Henze, G.P., Plamp, S., Cushing, V.J., 2013. Evaluation of commercial building HVAC systems as frequency regulation providers. *Energy and Buildings* 67, 225–235.
- Zhou, Q., Wang, S., Xu, X., Xiao, F., 2008. A grey-box model of next-day building thermal load prediction for energy-efficient control. *International Journal of Energy Research* 32, 1418–1431.
- Zong, Y., Su, W., Wang, J., Rodek, J.K., Jiang, C., Christensen, M.H., You, S., Zhou, Y., Mu, S., 2019. Model predictive control for smart buildings to provide the demand side flexibility in the multi-carrier energy context: Current status, pros and cons, feasibility and barriers. *Energy Procedia* 158, 3026–3031.
- Zur, R.M., Jiang, Y., Pesce, L.L., Drukker, K., 2009. Noise injection for training artificial neural networks: A comparison with weight decay and early stopping. *Medical Physics* 36, 4810–4818.



# Curriculum Vitae

**Felix Bünning**

**Born on September 8<sup>th</sup>, 1990 in Aachen, Germany**

2018 – 2021 **Doctoral Studies**

Automatic Control Laboratory,  
Department of Information Technology and Electrical Engineering,  
Swiss Federal Institute of Technology Zurich (ETH Zurich), Switzerland  
and  
Urban Energy Systems Laboratory,  
Department of Engineering Sciences,  
Swiss Federal Laboratories for Materials Science and Technology (Empa),  
Switzerland

2016 – 2017 **Research Affiliate**

Lawrence Berkeley National Laboratory,  
University of California, USA

2015 – 2017 **Master Studies**

Faculty of Mechanical Engineering,  
RWTH Aachen University, Germany

2011 – 2015 **Bachelor Studies**

Faculty of Mechanical Engineering,  
RWTH Aachen University, Germany

2010 – 2011 **Alternative civilian service**

Mulberry Bush School, England

2001 – 2010 **High School, German Abitur**

Einhard Gymnasium Aachen, Germany