

Microgrid management with weather-based forecasting of energy generation, consumption and prices.



Jonathan Dumas

Supervisor: Prof. Bertrand Cornélusse

Faculty of Applied Sciences
Department of Computer Science and Electrical Engineering
Liège University

This version of the manuscript is pending the approval of the jury.

This dissertation is submitted for the degree of
Doctor of Philosophy

Jury members

Prof. Quentin Louveaux (President) University of Liège, Belgium.

Prof. Bertrand Cornélusse (Supervisor) University of Liège, Belgium.

Prof. Gilles Louppe University of Liège, Belgium.

Prof. Pierre Pinson Technical University of Denmark, Denmark.

Prof. Ricardo Bessa INESC TEC, Portugal.

Prof. Simone Paoletti University of Siena, Italy.

“ Nothing is impossible, just do it and let’s see. ”

— **A true coconut-shaker**

“There are admirable potentialities in every human being. Believe in your strength and your youth. Learn to repeat endlessly to yourself: It all depends on me. ”

— **André Gide**

“In vino veritas. ”

— **Pliny the Elder, 23–79 A.D.**

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements.

Jonathan Dumas
November 2021

Acknowledgements

Throughout the writing of this dissertation, I have received a great deal of support and assistance.

Family

I would like to acknowledge my family, and especially Perrine and my daughter Lily. I could not have achieved this thesis without their emotional and material support. They always have been there for me. Perrine saved me much time taking care of our beloved daughter, **she did, in fact, the most challenging part of the job**. In my opinion, it is simple to work, especially in a research environment. It is much harder to raise a child. In addition, I could not have been so implied, without her support, in non-profit organizations to help make energy transition to a carbon-free economy: the Shifters Belgium¹, the working group on the sustainability of the ReD² (an association aiming to gather the Ph.D. candidates from the University of Liège), the climate³ and digital collages⁴, and the ClimActes 2021 summer school⁵.

Supervisor

I express my gratitude to my supervisor Prof. Bertrand Cornélusse, whose expertise was invaluable in formulating the research questions and methodology. Bertrand is significantly implicated in designing exciting lectures for bachelor and master students. He is also committed to supervising his Ph.D. students despite its busy agenda. Professors are supposed to be supermen or superwomen: teach, teach and teach, then, being at the front end of the research, and of course, implicated in the University and Faculty challenges. However, they have as everybody 24 hours per day :). He made it possible to work with the Sienna group and encouraged me throughout the past

¹<https://theshiftproject.org/equipe/#benevoles>

²<https://www.facebook.com/reduliege/>

³<https://climatefresk.org/>

⁴<https://digitalcollage.org/>

⁵<http://climactes.org/>

three years. We have had helpful discussions on technical topics like optimization and machine learning and friendly talks on many other questions. Your insightful feedback pushed me to sharpen my thinking and brought my work to a higher level. We had the opportunity to spend an excellent stay at the first microgrids summer school at Belfort in 2019. He also provided the funds for three summer schools: microgrids summer school at Belfort [2019], sustainable ICT summer school at Catholic University of Louvain [2020], and ClimActes at Liège University [2021]. He provided the funds for attending several conferences: 16th European Energy Market Conference (EEM 2019), XXI Power Systems Computation Conference (PSCC 2020), 16th International Conference on Probabilistic Methods Applied to Power Systems (PMAPS 2020), 14th IEEE PowerTech (PowerTech 2021). Finally, the two stays at Sienna University in 2019 and 2020 (just before the lockdown of the covid-19 crisis).

Thesis committee

I am grateful to Prof. Simone Paoletti, Prof. Quentin Louveaux, and Xavier Fettweis, research associate FNRS, for being part of the thesis committee, their support, and valuable guidance throughout this thesis and the discussions about optimization and weather forecasting. In addition, Xavier Fettweis taught how to use the MAR regional climate model. I want to thank all three of you for your regular support and the opportunities to further my research.

Thesis jury

I am deeply honored by this jury. Each one of the members is profoundly remarkable in his field and inspired my research during this thesis. I could not have dreamt of a better jury and hope to be at the level expected. In addition, I sincerely thank the jury members for devoting themselves to reading this manuscript and all the related administrative duties.

Liège University Applied Science Faculty colleagues

I thank my colleagues and friends of the Applied Science Faculty for their support and for the exciting discussions we had together during the last three years: especially Selmane Dakir, Antonio Sutera, Antoine Wehenkel, Laurine Duchesne, Pascal Leroy, Adrien Bolland, Raphael Fonteneau, Sylvain Quoilin, Gauthier Gain, Ioannis Noukas, Miguel Manuel de Villena, and all the others that I did not mention. I also thank Gilles Louppe, associate professor at Liège University, for the quality of his course "INFO8010 - Deep Learning". Gilles Louppe is very involved in designing exciting

courses. I enjoyed the lectures that I followed in Spring 2021 at the end of this thesis. I wish I had attended them at the beginning of my thesis to understand better the deep learning topic that I learned from scratch.

Coconut-shakers group and ReD

I thank my beloved friends of the *coconut-shakers group*: Antoine Dubois, François Rigo, Thibaut Theate, Gilles Corman, and Adrien Corman. We have worked together intensively to propose actions and ideas at the Applied Science Faculty to make tiny changes about the climate change issue. There is still a massive gap between rhetoric and actions. Nevertheless, we keep on believing that one-day things will change! I also thank the active members of the ReD: Kathleen Jacquerie, Chloe Stevenne, Coline Grégoire, and all the others that I did not mention.

Sienna University colleagues

I appreciated working with Prof. Simone Paoletti, Prof. Antonello Giannitrapani, and Prof. Antonio Vicino from Sienna University. It was terrific to work with them at Sienna University, and we have had enjoyable moments. They are devoted to their work, and they know how to make a good working relationship. They helped to understand stochastic programming better and to develop exciting research directions.

Master students

I thank all the master students who did incredible work on various topics such as forecasting, optimization during their master thesis, especially during the covid-19 crisis. It was a great pleasure to work with them: Audrey Lempereur [2020-Applied Science Faculty], Clement Liu [2020-Ecole Polytechnique], Jean Cyrus de Gourcuff [2021-Ecole Polytechnique], Colin Cointe [2021-Mines ParisTech], Damien Lanaspeze [2021-Mines ParisTech], Benjamin Delvoye [2021-Applied Science Faculty], and Gabriel Guerra Martin [2021-Polytechnic University of Catalonia].

Liège University Applied Science Faculty staff

I sincerely thank the shadow workers of the Applied Science Faculty: the cleaning workers (thank you Patricia!), the administrative staff (Eric Vangenechten), and all the others that I did not mention. I also thank the academic staff of the Applied Science Faculty: Eric Delhez (for the discussions on sustainability even if I think can do better!), Angélique Leonard, Dominique Toye, and all the others that I did not mention.

Liège University human resource staff

I thank all the human resource staff. We are so fortunate at the University of Liège to have free and very high-quality training courses on various topics such as self-development, negotiation, teamwork. In particular, I appreciate the training courses of Jean Yves Girin. He is incredible, and he helped me improve my human skills (but there is still much work to do). I can only encourage all the academics (including professors), administrative staff, Ph.D. candidates, and researchers to attend these training courses. I was the only man and the only one from the Applied Science Faculty most of the time. Does that mean that the 'men' do not need these training courses about human skills?

Liège University IFRES

I thank the IFRES for the pedagogical training that we attend as teaching assistants. We are lucky that Liège University provides these training courses.

P. Pinson

I thank professor Prof. Pierre Pinson from the Technical University of Denmark. His work has been a great inspiration during this thesis. In addition, he is open to discussion and sharing all material from his courses. In particular, I used a lot of the course materials "Renewables in Electricity Markets" to design the forecasting lessons that I taught in the course "ELEN0445 Microgrids" at the University of Liège during the years 2018, 2019, 2020, and 2021.

Summer schools

Microgrids Belfort 2019 summer school. I thank all the members of the microgrids Belfort 2019 summer school, and especially Robin Roche. It was the first summer school on microgrids, and we learned a lot in an enjoyable working environment. The Typhoon Hil team did a great job presenting its tools and organize the challenge.

Sustainable ICT 2020 summer school. I also thank the members of the Sustainable ICT 2020 summer school that took place at the Catholic University of Louvain (UCL). It was an excellent summer school with very skilled researchers and Ph.D. candidates. There is at UCL a group of Ph.D. candidates (Thibault Pirson, Gregoire Lebrun, and many others) and professors (Jean-Pierre Raskin, Hervé Jeanmart, David Bol) very motivated to address the sustainability challenges. It is a great source of inspiration.

ClimACTES 2021 summer school. I am grateful to all the members of the ClimACTES 2021 summer school that took place at Liège University. This event was outstanding where we have faced the challenges of climate change and the urgent need to reduce the emission of greenhouse gases drastically. The first week was dedicated to lessons from professors from several universities about these challenges. Then, we worked in the second week to develop projects to achieve ambitious targets to help to decarbonize the economy. I was part of the education team, and I learned from students and people from various backgrounds. It was one of the richest experiences in my life. Philippe Gilson, that founded ClimACTES, is a great source of inspiration. The work of all the ClimACTES volunteers was incredible. They dedicated two entire years to build this summer school. I can only say congratulations!

Reviewers

I would like to thank all the reviewers of the papers submitted to conferences and journals. They helped me to improve the work proposed and to get out of my comfort zone.

Liège University

Finally, I thank the University of Liège. Indeed, they financed me as a teaching assistant during these two last years and provided me with a nice working environment.

Abstract

“*Do. Or do not. There is no try.*”

— Yoda

Climate Change

The decade 2009-2019 was particularly intense in rhetoric about efforts to tackle the climate crisis, such as the 2015 United Nations Climate Change Conference, COP 21. However, the carbon dioxide emissions at the world scale increased constantly from 29.7 (GtCO₂) in 2009 to 34.2 in 2019. *The current gap between rhetoric and reality on emissions was and is still huge.* The Intergovernmental Panel on Climate Change proposed different mitigation strategies to achieve the net emissions reductions that would be required to follow a pathway that **limits global warming to 1.5°C** with no or limited overshoot. There are still pathways to reach net-zero by 2050. Several reports propose detailed scenarios and strategies to achieve these targets. They remain narrow and highly challenging, requiring all stakeholders, governments, businesses, investors, and citizens to take action this year and every year after so that the goal does not slip out of reach. In most of these trajectories, *electrification* and an increased share of *renewables* are some of the key pillars. The transition towards a carbon-free society goes through **an inevitable increase in the share of renewable generation in the energy mix** and a drastic decrease in the total consumption of fossil fuels.

Thesis topic

In contrast to conventional power plants, renewable energy is subject to uncertainty. Most of the generation technologies based on renewable sources are non-dispatchable, and their production is stochastic and complex to predict in advance. A high share of renewables is challenging for power systems that have been designed and sized for dispatchable units. Therefore, this thesis studies the *integration of renewables in power systems* by investigating forecasting and decision-making tools.

Since variable generation and electricity demand both fluctuate, they must be forecast ahead of time to inform real-time electricity scheduling and longer-term system

planning. Better short-term forecasts enable system operators to reduce reliance on polluting standby plants and proactively manage increasing amounts of variable sources. Better long-term forecasts help system operators and investors to decide where and when to build variable plants. In this context, *probabilistic* forecasts, which aim at modeling the distribution of all possible future realizations, have become a vital tool to equip decision-makers, hopefully leading to better decisions in energy applications.

When balancing electricity systems, system operators use scheduling and dispatch to determine how much power every controllable generator should produce. This process is slow and complex, governed by NP-hard optimization problems such as unit commitment and optimal power flow. Scheduling becomes even more complex as electricity systems include more storage, variable generators, and flexible demand. Thus, scheduling must improve significantly, allowing operators to rely on variable sources to manage systems. Therefore, stochastic or robust optimization strategies have been developed along with decomposition techniques to make the optimization problems tractable and efficient.

Thesis content

These two challenges raise two central research questions studied in this thesis: (1) How to produce reliable probabilistic renewable generation forecasts, consumption, and electricity prices? (2) How to make decisions with uncertainty using probabilistic forecasts to improve scheduling? The thesis perimeter is the energy management of "small" systems such as *microgrids* at a residential scale on a day-ahead basis. The manuscript is divided into two main parts to propose directions to address both research questions: (1) a forecasting part; (2) a planning and control part.

Thesis first part

The forecasting part presents several techniques and strategies to produce and evaluate probabilistic forecasts. We provide the forecasting basics by introducing the different types of forecasts to characterize the behavior of stochastic variables, such as renewable generation, and the tools to assess the different types of forecasts. An example of forecast quality evaluation is given by assessing PV and electrical consumption point forecasts computed by standard deep-learning models such as recurrent neural networks. Then, the following Chapters investigate the quantile forecasts, scenarios, and density forecasts on several case studies. First, more advanced deep-learning models such as the encoder-decoder architecture produce PV quantile forecasts. Second, a density forecast-based approach computes probabilistic forecasts of imbalance prices

on the Belgian case. Finally, a recent class of deep generative models, normalizing flows, generates renewable production and electrical consumption scenarios. Using an energy retailer case study, this technique is extensively compared to state-of-the-art generative models, the variational autoencoders and generative adversarial networks.

Thesis second part

The planning and control part proposes approaches and methodologies based on optimization for decision-making under uncertainty using probabilistic forecasts on several case studies. We introduce the basics of decision-making under uncertainty using optimization strategies: stochastic programming and robust optimization. Then, we investigate these strategies in several case studies in the following Chapters. First, we propose a value function-based approach to propagate information from operational planning to real-time optimization in a deterministic framework. Second, three Chapters focus on the energy management of a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. This framework promotes renewable power generation facilities in small non-interconnected grids. The day-ahead planning of the system uses either a stochastic or a robust approach. Furthermore, a sizing methodology of the system is proposed. Finally, we consider the day-ahead market scheduling of an energy retailer to evaluate the forecast value of the deep learning generative models introduced in the forecasting part.

Perspectives

We propose four primary research future directions. (1) Forecasting techniques of the future. The development of new machine learning models that take advantage of the underlying physical process opens a new way of research. For instance, new forecasting techniques that take advantage of the power system characteristics, such as the graphical normalizing flows capable of learning the power network structure, could be applied to hierarchical forecasting. (2) Machine learning for optimization. Models that simplify optimization planning problems by learning a sub-optimal space. For instance, a deep learning model can partially learn the sizing space to provide a fast and efficient tool. A neural network can also emulate the behavior of a physics solver that solves electricity differential equations to compute electricity flow in power grids. Furthermore, such proxies could evaluate if a given operation planning decision would lead to acceptable trajectories where the reliability criterion is met in real-time. (3) Modelling and simulation of energy systems. New flexible and open-source optimization modeling tools are required to capture the growing complexity of such

future energy systems. To this end, in the past few years, several open-source models for the strategic energy planning of urban and regional energy systems have been developed. EnergyScope TD and E4CLIM are two of them where we think it may be relevant to implement and test the forecasting techniques and scheduling strategies developed in this thesis. (4) Psychology and machine learning. Achieving sustainability goals requires as much the use of relevant technology as psychology. Therefore, one of the main challenges is not designing relevant technological tools but changing how we consume and behave in our society. Thus, machine learning and psychology could help to identify appropriate behaviors to reduce carbon footprint. Then, inform individuals, and provide constructive opportunities by modeling individual behavior.

Table of contents

Foreword	1
1 General introduction	9
1.1 Context and motivations	9
1.2 Classification of forecasting studies	13
1.3 Content and contributions	14
I Forecasting	21
Part I general nomenclature	25
2 Forecasting background	27
2.1 Point forecast	29
2.2 Probabilistic forecasts	30
2.2.1 Quantiles	30
2.2.2 Prediction intervals	31
2.2.3 Scenarios	32
2.2.4 Density forecasts	32
2.3 Model-based formulation	32
2.4 Model training	33
2.4.1 Regression with supervised learning	33
2.4.2 Empirical risk minimization	34
2.4.3 The "double descent" curve	36
2.4.4 Training methodology	37
2.4.5 Learning, validation, and testing sets	37
2.4.6 k-fold cross-validation	38
2.5 Conclusions	38

3	Forecast evaluation	41
3.1	Metrics for point forecasts	43
3.2	Metrics for probabilistic forecasts	44
3.2.1	Calibration	44
3.2.2	Univariate skill scores	45
3.2.3	Multivariate skill scores	47
3.3	Point forecasting evaluation example	49
3.3.1	Formulation	49
3.3.2	Forecasting models	50
3.3.3	Results	51
3.4	Conclusions	53
4	Quantile forecasting	55
4.1	Formulation	58
4.2	Related work	59
4.3	Forecasting models	59
4.3.1	Gradient boosting regression (GBR)	59
4.3.2	Multi-layer perceptron (MLP)	60
4.3.3	Encoder-decoder (ED)	60
4.4	The ULiège case study	61
4.4.1	Case study description	61
4.4.2	Numerical settings	62
4.4.3	Day-ahead results	63
4.4.4	Intraday results	64
4.4.5	Conclusions	68
4.5	Comparison with generative models	68
4.6	Conclusions	70
5	Density forecasting of imbalance prices	71
5.1	Related work	73
5.2	Formulation	74
5.2.1	Net regulation volume forecasting	74
5.2.2	Imbalance price forecasting	75
5.3	Case study	77
5.4	Results	78
5.5	Conclusions	81
5.6	Appendix: balancing mechanisms	83

5.6.1	Balancing mechanisms	83
5.6.2	Belgium balancing mechanisms	84
5.6.3	Additional results	84
6	Scenarios	87
6.1	Introduction	90
6.1.1	Related work	91
6.1.2	Research gaps and scientific contributions	93
6.1.3	Applicability of the generative models	97
6.1.4	Organization	97
6.2	Background	98
6.2.1	Multi-output forecasts using a generative model	98
6.2.2	Deep generative models	98
6.2.3	Theoretical comparison	105
6.3	Quality assessment	106
6.3.1	Classifier-based metric	107
6.3.2	Correlation matrix between scenarios	109
6.3.3	Diebold-Mariano test	109
6.4	Case study	111
6.4.1	Implementation details	111
6.4.2	Hyper-parameters	114
6.5	Quality results	117
6.5.1	Wind track	117
6.5.2	All tracks	118
6.6	Conclusions and perspectives	125
6.7	Appendix: Table 6.1 justifications	126
6.8	Appendix: background	127
6.8.1	NFs	127
6.8.2	VAEs	128
6.8.3	GANs	130
7	Part I conclusions	133
II	Planning and control	135
8	Decision-making background	139
8.1	Linear programming	141

8.1.1	Formulation of a linear programming problem	141
8.1.2	Duality in linear programming	141
8.2	Stochastic optimization	142
8.3	Robust optimization	144
8.3.1	Benders-dual cutting plane algorithm	145
8.3.2	Column and constraints generation algorithm	146
8.4	Conclusions	147
9	Coordination of the planner and controller	149
9.1	Notation	153
9.2	Problem statement	155
9.2.1	Assumptions	155
9.2.2	Formulation	155
9.3	Proposed method	156
9.3.1	Computing the cost-to-go function	157
9.3.2	OPP formulation	158
9.3.3	OP constraints	159
9.3.4	RTP formulation	160
9.3.5	RTO constraints	161
9.4	Test description	161
9.5	Numerical results	163
9.5.1	No symmetric reserve	163
9.5.2	Results with symmetric reserve	165
9.6	Conclusions	166
10	Capacity firming using a stochastic approach	169
10.1	Notation	172
10.1.1	Sets and indices	172
10.1.2	Parameters	173
10.1.3	Variables	174
10.2	The Capacity Firming Framework	174
10.2.1	Day-ahead engagement	175
10.2.2	Real-time control	175
10.3	Problem formulation	176
10.3.1	Deterministic approach	176
10.3.2	Deterministic approach with perfect forecasts	178
10.3.3	Stochastic approach	179

10.3.4	Evaluation methodology	179
10.4	MiRIS microgrid case study	180
10.4.1	Results for unbiased PV scenarios with fixed variance	182
10.4.2	BESS capacity sensitivity analysis	184
10.5	Conclusions and perspectives	187
10.6	Appendix: PV scenario generation	187
11	Capacity firming sizing	191
11.1	Problem statement	194
11.1.1	Stochastic approach	195
11.1.2	Deterministic approach	196
11.1.3	Oracle controller	196
11.2	Forecasting methodology	196
11.2.1	Gaussian copula-based PV scenarios	197
11.2.2	PV point forecast parametric model	197
11.2.3	PV scenarios	198
11.3	Sizing study	198
11.3.1	Problem statement and assumptions	198
11.3.2	Levelized cost of energy (LCOE)	200
11.3.3	Case study description	201
11.3.4	Sizing parameters	202
11.3.5	Sizing results	202
11.4	Conclusions and perspectives	204
12	Capacity firming using a robust approach	207
12.1	Problem formulation	212
12.1.1	Deterministic planner formulation	213
12.1.2	Robust planner formulation	213
12.1.3	Second-stage planner transformation	215
12.1.4	Controller formulation	217
12.2	Solution methodology	217
12.2.1	Convergence warm start	218
12.2.2	Algorithm convergence	218
12.2.3	Benders-dual cutting plane algorithm	219
12.2.4	Column and constraints generation algorithm	221
12.3	Case Study	222
12.3.1	Numerical settings	223

12.3.2	Constant risk-averse parameters strategy	223
12.3.3	Dynamic risk-averse parameters strategy	225
12.3.4	BD convergence warm start improvement	228
12.3.5	BD and CCG comparison	229
12.4	Conclusion	230
13	Energy retailer	231
13.1	Notation	233
13.2	Problem formulation	234
13.2.1	Stochastic planner	234
13.2.2	Dispatching	236
13.3	Value results	237
13.3.1	Results summary	239
13.4	Conclusions	239
13.5	Appendix: Table 13.2 justifications	241
14	Part II conclusions	243
15	General conclusions and perspectives	245
15.1	Summary	245
15.2	Future directions	248
15.2.1	Forecasting techniques of the future	248
15.2.2	Machine learning for optimization	248
15.2.3	Modelling and simulation of energy systems	250
15.2.4	Machine learning and psychology	251
	References	253

Foreword

“The greatest glory in living lies not in never falling, but in rising every time we fall.”

— **Nelson Mandela**

“Adults keep saying we owe it to the young people, to give them hope, but I don't want your hope. I don't want you to be hopeful. I want you to panic. I want you to feel the fear I feel every day. I want you to act. I want you to act as you would in a crisis. I want you to act as if the house is on fire, because it is.”

— **Greta Thunberg**

Climate Change

The Intergovernmental Panel on Climate Change (IPCC) set the global net anthropogenic CO₂ emission trajectories and targets, depicted in Figure 1, to limit climate change [6]. It is a summary for policymakers (SPM) that presents the key findings of the special report published in 2018⁶, based on the assessment of the available scientific, technical and socio-economic literature relevant to Global Warming of 1.5°C and for the comparison between Global Warming of 1.5°C and 2°C above pre-industrial levels. The SPM presents the emission scenarios consistent with 1.5°C Global Warming:

In model pathways with no or limited overshoot of 1.5°C, global net anthropogenic CO₂ emissions decline by about 45% from 2010 levels by 2030

⁶The IPCC released in August 2021 the AR6 SPM [42]. It presents key findings of the Working Group I contribution to the IPCC's Sixth Assessment Report (AR6) on the physical science basis of climate change. It is a must to read for every decision-maker, professor, researcher, or person that wants to understand the challenges at stake.


(40–60% interquartile range), reaching net-zero around 2050 (2045–2055 interquartile range).⁷ [6][C.1]⁷

Global emissions pathway characteristics

General characteristics of the evolution of anthropogenic net emissions of CO₂, and total emissions of methane, black carbon, and nitrous oxide in model pathways that limit global warming to 1.5°C with no or limited overshoot. Net emissions are defined as anthropogenic emissions reduced by anthropogenic removals. Reductions in net emissions can be achieved through different portfolios of mitigation measures illustrated in Figure SPM.3b.

Global total net CO₂ emissions

Billion tonnes of CO₂/yr




Timing of net zero CO₂
Line widths depict the 5-95th percentile and the 25-75th percentile of scenarios




Non-CO₂ emissions relative to 2010

Emissions of non-CO₂ forcers are also reduced or limited in pathways limiting global warming to 1.5°C with **no or limited overshoot**, but they do not reach zero globally.


Methane emissions



Black carbon emissions



Nitrous oxide emissions



Source: IPCC Special Report on Global Warming of 1.5°C

Fig. 1 IPCC global net anthropogenic CO₂ emission pathways and targets. Credits: [6, Figure SPM.3a].

The IPCC proposes different mitigation strategies to achieve the net emissions reductions required to follow a pathway to limit Global Warming to 1.5°C with no or limited overshoot. These strategies require a tremendous amount of effort from all countries.

⁷<https://www.ipcc.ch/sr15/chapter/spm/>

'Pathways limiting Global Warming to 1.5°C with no or limited overshoot would require rapid and far-reaching transitions in energy, land, urban and infrastructure (including transport and buildings), and industrial systems (high confidence). These systems transitions are unprecedented in terms of scale, but not necessarily in terms of speed, and imply deep emissions reductions in all sectors, a wide portfolio of mitigation options and a significant upscaling of investments in those options (medium confidence).'
[6][C.2]⁷

Trajectories to achieve IPCC targets

Several reports propose pathways and strategies to achieve these targets. The International Energy Agency (IEA⁸) presents a comprehensive study of how to transition to a net-zero energy system by 2050 in the special report 'Net-zero by 2050: A roadmap for the global energy system' [3]. It is consistent with limiting the global temperature rise to 1.5 °C without a temperature overshoot (with a 50 % probability). The Net-Zero Emissions by 2050 Scenario (NZE) shows what is needed for the global energy sector to achieve net-zero CO₂ emissions by 2050:

'In the NZE, global energy-related and industrial process CO₂ emissions fall by nearly 40% between 2020 and 2030 and to net-zero in 2050. Universal access to sustainable energy is achieved by 2030. There is a 75% reduction in methane emissions from fossil fuel use by 2030. These changes take place while the global economy more than doubles through to 2050 and the global population increases by 2 billion.' [3, Chapter 2: Summary]

The key pillars of decarbonization of the global energy system proposed are (1) energy efficiency, (2) behavioral changes, (3) **electrification**, (4) **renewables**, (5) hydrogen and hydrogen-based fuels, (6) bioenergy, and (7) carbon capture, utilization, and storage.

First, concerning electrification, the direct use of low-emissions electricity in place of fossil fuels is one of the most significant drivers of emissions reductions in the NZE, accounting for around 20% of the total reduction achieved by 2050. The share of the electricity in the final consumption increases from 20% in 2020 to 49% in 2050. It concerns the sectors of the industry, transport, and buildings. Second, concerning renewables:

'At a global level, renewable energy technologies are the key to reducing emissions from electricity supply. Hydropower has been a leading low-

⁸<https://www.iea.org/>

emission source for many decades, but it is mainly the expansion of wind and solar that triples renewables generation by 2030 and increases it more than eightfold by 2050 in the NZE. The share of renewables in total electricity generation globally increases from 29% in 2020 to over 60% in 2030 and to nearly 90% in 2050. To achieve this, annual capacity additions of wind and solar between 2020 and 2050 are five-times higher than the average over the last three years. Dispatchable renewables are critical to maintain electricity security, together with other low-carbon generation, energy storage and robust electricity networks. In the NZE, the main dispatchable renewables globally in 2050 are hydropower (12% of generation), bioenergy (5%), concentrating solar power (2%) and geothermal (1%).’ [3, Section 2.4.5]

The IEA report is not the ground truth but has the merit to propose guidelines and directions. There are many other reports and organizations that present strategies and scenarios to achieve the IPCC targets. For instance, The Shift Project (TSP) is a European think tank⁹ advocating the shift to a post-carbon economy. It proposes guidelines and information on energy transition in Europe. The take-home message is that **there are still pathways to reach net-zero by 2050**. They remain narrow and challenging, requiring all stakeholders, governments, businesses, investors, and citizens to take action this year and every year after so that the goal does not slip out of reach.

Gap between rhetoric and reality

However, the **current gap between rhetoric and reality on emissions is still huge**:

’We are approaching a decisive moment for international efforts to tackle the climate crisis – a great challenge of our times. The number of countries that have pledged to reach net-zero emissions by mid-century or soon after continues to grow, but so do global greenhouse gas emissions. This gap between rhetoric and action needs to close if we are to have a fighting chance of reaching net-zero by 2050 and limiting the rise in global temperatures to 1.5 °C.’ [3, Foreword]

Figure 2 illustrates humorously this gap. The UNEP Emissions Gap Report provides every year a review of the difference between the greenhouse emissions forecast in 2030 and where they should be to avoid the worst impacts of climate change. Figure

⁹<https://theshiftproject.org/en/home/>




Fig. 2 Climate change fake news. Credits: Xavier Gorce.

3 depicts the global GHG emissions under different scenarios and the emissions gap in 2030. 2009-2019 was particularly intense in rhetoric tackling the climate crisis like the 2015 United Nations Climate Change Conference, COP 21. However, the carbon dioxide emissions at the world scale constantly rose from 29.7 (GtCO₂) in 2009 to 34.2 in 2019 [118]¹⁰. In the meantime, the primary energy consumption increased from 134 000 TWh to 162 000. In 2019, the primary energy consumption by fuel was composed of oil 53 600 (33.1%), coal 39 300 (24.2%), natural gas 43 900 (27.0%), nuclear energy 6 900 (4.3%), hydro-electricity 10 500 (6%), and renewables 8 100 (5%), as depicted in Figure 4. The share of renewables in the energy mix progressed and reached 5% in 2019 (a record). However, the total consumption of fossil fuels such as oil, coal, and natural gas also rose. Oil continues to hold the largest share of the energy mix, coal is the second-largest fuel, and natural gas grew to a record share of 24.2%.

The Covid-19 pandemic has delivered a shock to the world economy. It resulted in an unprecedented 5.8% decline in CO₂ emissions in 2020. The IPCC targets for 2050 require a reduction of 5% each year from 2020 to 2050. However, the IEA data shows that global energy-related CO₂ emissions have started to climb again since December 2020. Nevertheless, there is still hope, and every action to decrease the CO₂ emissions to gain a reduction of 0.1°C is a winning!

¹⁰<https://www.bp.com/content/dam/bp/business-sites/en/global/corporate/pdfs/energy-economics/statistical-review/bp-stats-review-2020-full-report.pdf>




Fig. 3 Global GHG emissions under different scenarios and the emissions gap in 2030 (median and 10-th to 90-th percentile range; based on the pre-COVID-19 current policies scenario). Credits: [130, Figure ES.5].




Fig. 4 World consumption of primary energy (left) and shares of global primary energy (right) from 1994 to 2020. Credits: BP's Statistical Review of World Energy 2020 [118].

Chapter 1

General introduction

Overview

This chapter introduces the context, motivations, content, and contributions of the thesis. Two main parts compose the manuscript: (1) forecasting; (2) planning and control. Finally, it lists the publications.

“*Life has no meaning a priori... It is up to you to give it a meaning, and value is nothing but the meaning that you choose.*”

— **Jean-Paul Sartre**

1.1 Context and motivations

Statement 1. *Let suppose a utopian world where the current gap between rhetoric and reality on GHG emissions has been drastically decreasing to limit Climate Change and achieve the ambitious targets prescribed by the IPCC.*

Therefore, according to the IPCC targets, the transition to a carbon-free society goes through an inevitable increase in the renewable generation’s share in the energy mix and a drastic decrease in the total consumption of fossil fuels.

Statement 2. *This thesis does not debate or study whether and where renewables should be implemented.*

Renewables have pros and cons and are not carbon-free. We take the scenarios proposed by organizations such as the IPCC or IEA that evaluate the relevance of a particular type of renewable energy in the energy transition pathways. In addition,

the development of variable renewable energies systems results from and produces many interactions between several actors such as citizens, electricity networks, markets, and ecosystems [109]. This thesis does not discuss energy transformations' political and social aspects and focuses on integrating renewable energies into an existing interconnected system.

Statement 3. *This thesis studies the integration of renewables in power systems by investigating forecasting and decision-making tools based on machine learning.*

In contrast to conventional power plants, renewable energy is subject to uncertainty. Generation technologies based on renewable sources, with the notable exception of hydro and biomass, are *non-dispatchable*, *i.e.*, their output cannot or can only partly be controlled at the will of the producer. Their production is *stochastic* [124] and therefore, hard to forecast. A high share of renewables is challenging for power systems that have been designed and sized for dispatchable units [168]. Variable renewable energies depend on meteorological conditions and thus pose challenges to the electricity system's adequacy when conventional capacities are reduced. Therefore, it is necessary to redefine the flexible power system features [90].

Machine learning can contribute on all fronts by informing the research, deployment, and operation of electricity system technologies. High leverage contributions in power systems include [151]: accelerating the development of clean energy technologies, *improving demand and clean energy forecasts*, *improving electricity system optimization and management*, and enhancing system monitoring. This thesis focuses on two leverages: (1) the supply and demand forecast; (2) the electricity system optimization and management.

Since variable generation and electricity demand both fluctuate, they must be forecast ahead of time to inform real-time electricity scheduling and longer-term system planning. Better short-term forecasts enable system operators to reduce reliance on polluting standby plants and proactively manage increasing amounts of variable sources. Better long-term forecasts help system operators and investors to decide where and when to build variable plants. Forecasts need to become more accurate, span multiple horizons in time and space, and better quantify uncertainty to support these use cases. In this context, *probabilistic* forecasts [69], which aim at modeling the distribution of all possible future realizations, have become an important tool to equip decision-makers, hopefully leading to better decisions in energy applications [124, 85, 87].

When balancing electricity systems, system operators use scheduling and dispatch to determine how much power every controllable generator should produce. This process is slow and complex, governed by NP-hard optimization problems [151] such

as unit commitment and optimal power flow that must be coordinated across multiple time scales, from sub-second to days ahead. Scheduling becomes even more complex as electricity systems include more storage, variable generators, and flexible demand. Indeed, operators manage even more system components while simultaneously solving scheduling problems more quickly to account for real-time variations in electricity production. Thus, scheduling must improve significantly, allowing operators to rely on variable sources to manage systems.

Therefore, the two main research questions are:

1. How to produce reliable probabilistic forecasts of renewable generation, consumption, and electricity prices?
2. How to make decisions with uncertainty using probabilistic forecasts to improve scheduling?

Modeling tools for energy systems differ in terms of their temporal and spatial resolutions, level of technical details, simulation horizons. In particular, two classes of models can be distinguished depending on their focus: (1) system operation and (2) system design [35]. [114, Chapter 1] depicts both of these classes.

In the first class, the operational models of energy systems generally optimize a single sector's operation, such as the electricity sector, and do not consider investment costs. They describe the constraints of the studied system with a high level of accuracy and can model rapid variations in renewable energy production, forecast errors, or reserve markets. The energy system must be controlled at any time, which requires an excellent technical resolution of the components and accounts for the risks of failure. The system must decide how the load is shared between the different units below the hour at a short time horizon. Solving this problem involves accurately representing production units, such as the power ramps up or load constraints. Therefore, models will decide which units must be committed to optimally dispatching the load or generating shortly (hours - days).

In the second class, the long-term planning models generate scenarios for an energy system's long-term evolution. They include investments, optimize the system design over multiple years, and have a lower technical resolution of operational constraints. With a horizon of one year, new units can be built in existing sites or former units modernized. With a longer horizon, the overall system can be changed. When the horizon is long enough to neglect the existing system, the models can then optimize the future design of the energy system from scratch and assess the impact of different long-term policies on the design of the system.

Each model class is essential and answers different needs to tackle the energy transition and help integrate renewable energies. This thesis addresses the first class of models. An example of a thesis investigating the second one is Limpens [114].

Statement 4. *This thesis considers the energy management of "small" systems such as microgrids at a residential scale on a day-ahead basis.*

Indeed, the development of microgrids provides an effective way to integrate renewable energy sources and exploit the available flexibility in a decentralized manner. Microgrids are small electrical networks composed of decentralized energy resources and loads controlled locally. They can be operated either interconnected or in islanded mode. Figure 1.1 depicts a microgrid composed of PV generation, diesel generator (Genset), storage systems, load, and an energy management system (EMS). Energy storage is a crucial component for the stable and safe operation of a microgrid. Storage devices can compensate for the variability of the renewable energy sources and the load to balance the system. The reader is referred to Zia et al. [191] that proposes a comparative and critical analysis on decision-making strategies and their solution methods for microgrid energy management systems.

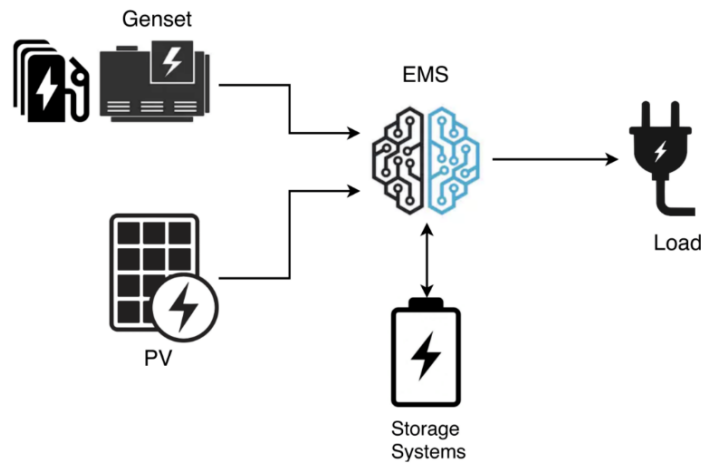


Fig. 1.1 Microgrid scheme. Credits: ELEN0445 Microgrids course <https://github.com/bcornelusse/ELEN0445-microgrids>, Liège University.

Statement 5. *This thesis considers the energy management of grid-connected microgrids on a day-ahead basis.*

Therefore, we are interested in producing reliable day-ahead probabilistic forecasts of renewable generation, consumption, and electricity prices¹ for a microgrid composed

¹This thesis considers only the imbalance prices in the Belgian case study.

of PV or wind generation and electrical consumption. However, in some specific cases, this perimeter is not strictly respected. For instance, the sizing of a grid-connected PV plant with a battery energy storage system is studied in the specific framework of capacity firming or the day-ahead planning of an energy retailer.

1.2 Classification of forecasting studies

One of the first works of this thesis was to conduct a literature review of forecasting studies. The forecasting literature is vast and composed of thousands of papers, even when selecting a particular field such as load or PV forecasting. Therefore, a *classification* into *two dimensions* of load forecasting studies is proposed by Dumas and Cornélusse [48] to decide which forecasting tools to use in which case. The approach can be extended to electricity prices, PV, or wind power forecasting. This classification aims to provide a synthetic view of the relevant forecasting techniques and methodologies by forecasting problem. This methodology is illustrated by reviewing several papers and summarizing the leading techniques and methodologies' fundamental principles.

The classification process relies on two parameters that define a forecasting problem: a temporal couple with the forecasting horizon and the resolution and a load couple with the system size and the load resolution. Each article is classified with key information about the dataset used and the forecasting tools implemented: the forecasting techniques (probabilistic or deterministic) and methodologies, the data cleansing techniques, and the error metrics. The process to select the articles reviewed was conducted into two steps. First, a set of load forecasting studies was built based on relevant load forecasting reviews and forecasting competitions. The second step consisted of selecting the most relevant studies of this set based on the following criteria: the quality of the description of the forecasting techniques and methodologies implemented, the description of the results and the contributions. For the sake of clarity, this manuscript does not detail this study. It can be read in two passes.

1. The first one identifies the forecasting problem of interest to select the corresponding class into one of the four classification tables. Each one references all the articles classified across a forecasting horizon. They provide a synthetic view of the forecasting tools used by articles addressing similar forecasting problems. Then, a second level composed of four Tables summarizes key information about the forecasting tools and the results of these studies.

- The second pass consists of reading the key principles of the main techniques and methodologies of interest and the reviews of the articles.

1.3 Content and contributions

The manuscript is divided into two main parts: Part I forecasting; Part II planning and control. Figure 1.2 depicts the thesis skeleton. Part I provides the forecasting tools and metrics required to produce and evaluate reliable point and probabilistic forecasts to be used as input of decision-making models in Part II. The latter proposes approaches and methodologies based on optimization for decision-making under uncertainty using probabilistic forecasts on several case studies.

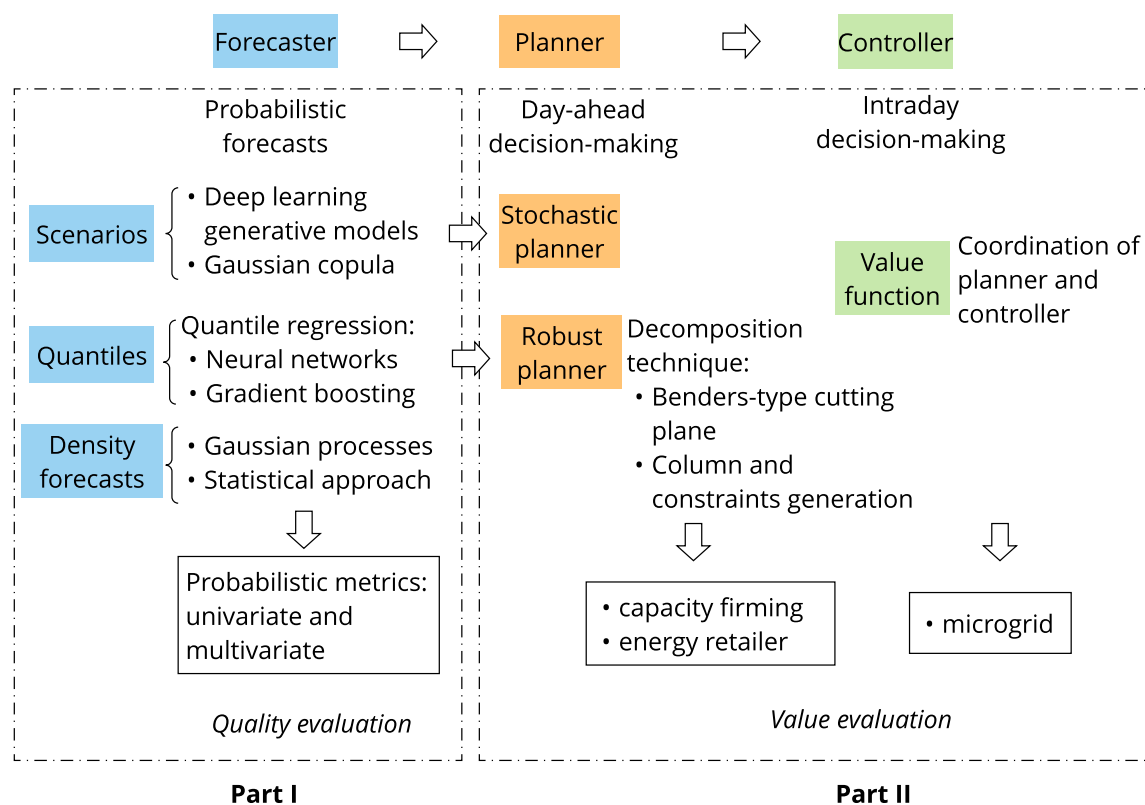


Fig. 1.2 Thesis skeleton.

Part I content and contributions

The content and main contributions of the forecasting part are:

- Chapter 2 introduces different types of forecasts to characterize the behavior of stochastic variables, such as renewable generation, electrical consumption, and electricity prices.
- Chapter 3 provides the tools to assess the different types of forecasts. For predictions in any form, one must differentiate between their quality and their value. This Chapter focus on forecast quality. Part II addresses the forecast value. An example of forecast quality assessment is conducted on PV and electrical consumption point forecasts. They are computed using common deep-learning models such as recurrent neural networks and used as inputs of day-ahead planning in Chapter 9.

References: The point forecast quality evaluation is an extract of Jonathan Dumas, Selmane Dakir, Clément Liu, and Bertrand Cornélusse. Coordination of operational planning and real-time optimization in microgrids. *Electric Power Systems Research*, 190:106634, 2021. URL <https://arxiv.org/abs/2106.02374>.

- Chapter 4 investigates PV quantiles forecasts using deep-learning models such as the encoder-decoder architecture. Then, Chapter 12 uses the PV intraday point and quantiles forecasts as inputs of a robust optimization planner.

References: This chapter is an adapted version of Jonathan Dumas, Colin Cointe, Xavier Fettweis, and Bertrand Cornélusse. Deep learning-based multi-output quantile forecasting of pv generation. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494976. URL <https://arxiv.org/abs/2106.01271>.

- Chapter 5 proposes imbalance prices density forecasts with a particular focus on the Belgian case. A two-step density forecast-based approach computes the net regulation volume state transition probabilities to infer the imbalance prices.

References: This chapter is an adapted version of Jonathan Dumas, Ioannis Boukas, Miguel Manuel de Villena, Sébastien Mathieu, and Bertrand Cornélusse. Probabilistic forecasting of imbalance prices in the belgian context. In *2019 16th International Conference on the European Energy Market (EEM)*, pages 1–7. IEEE, 2019. URL <https://arxiv.org/abs/2106.07361>.

- Chapter 6 studies the generation of scenarios for renewable production and electrical consumption by implementing a recent class of deep generative models, normalizing flows. It provides a fair comparison of the quality and value of this technique with the state-of-the-art deep learning generative models, Variational AutoEncoders, and Generative Adversarial Networks. Chapter 13 assesses the forecast value.

References: This chapter is an adapted version of Jonathan Dumas, Antoine Wehenkel, Damien Lanaspèze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

- Finally, Chapter 7 draws the general conclusions and perspectives of Part I.

Part II content and contributions

The content and main contributions of the planning and control part are:

- Chapter 8 introduces different types of optimization strategies for decision-making under uncertainty: stochastic and robust optimization. Then, it presents succinctly two decomposition methods to address the two-stage robust non-linear optimization problem: the Benders dual-cutting plane and the column and constraints generation algorithms.
- Chapter 9 presents a value function-based approach as a way to propagate information from operational planning to real-time optimization.

References: This chapter is an adapted version of Jonathan Dumas, Selmane Dakir, Clément Liu, and Bertrand Cornélusse. Coordination of operational planning and real-time optimization in microgrids. *Electric Power Systems Research*, 190:106634, 2021. URL <https://arxiv.org/abs/2106.02374>.

- Chapter 10 addresses the energy management, using a stochastic approach, of a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. This framework has been designed to promote renewable power generation facilities in small non-interconnected grids.

References: This chapter is an adapted version of Jonathan Dumas, Bertrand Cornélusse, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino.

Stochastic and deterministic formulations for capacity firming nominations. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–7. IEEE, 2020. URL <https://arxiv.org/abs/2106.02425>.

- Chapter 11 extends Chapter 10 and proposes a sizing methodology of the system.

References: This chapter is an adapted version of Jonathan Dumas, Bertrand Cornélusse, Xavier Fettweis, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Probabilistic forecasting for sizing in the capacity firming framework. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494947. URL <https://arxiv.org/abs/2106.02323>.

- Chapter 12 extends Chapter 10 and investigates the day-ahead planning using robust optimization.

Python code: <https://github.com/jonathandumas/capacity-firming-ro>

References: This chapter is an adapted version of Jonathan Dumas, Colin Cointe, Antoine Wehenkel, Antonio Sutera, Xavier Fettweis, and Bertrand Cornélusse. A probabilistic forecast-driven strategy for a risk-aware participation in the capacity firming market. *IEEE Transactions on Sustainable Energy*, pages 1–1, 2021. doi: 10.1109/TSTE.2021.3117594. URL <https://arxiv.org/abs/2105.13801>.

- Chapter 13, is the extension of Chapter 6, and presents the forecast value evaluation of the deep learning generative models by considering the day-ahead market scheduling of electricity aggregators, such as energy retailers or generation companies.

Python code: <https://github.com/jonathandumas/generative-models>

References: This chapter is an adapted version of Jonathan Dumas, Antoine Wehenkel, Damien Lanaspèze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

- Finally, Chapter 14 draws the general conclusions and perspectives of Part II.

Publications

The thesis is mainly based on the following studies, all available in open-access on arXiv, listed in chronological order:

- Jonathan Dumas and Bertrand Cornélusse. Classification of load forecasting studies by forecasting problem to select load forecasting techniques and methodologies, 2020. URL <https://arxiv.org/abs/1901.05052>
- Jonathan Dumas, Ioannis Boukas, Miguel Manuel de Villena, Sébastien Mathieu, and Bertrand Cornélusse. Probabilistic forecasting of imbalance prices in the belgian context. In *2019 16th International Conference on the European Energy Market (EEM)*, pages 1–7. IEEE, 2019. URL <https://arxiv.org/abs/2106.07361>
- Jonathan Dumas, Selmane Dakir, Clément Liu, and Bertrand Cornélusse. Coordination of operational planning and real-time optimization in microgrids. *Electric Power Systems Research*, 190:106634, 2021. URL <https://arxiv.org/abs/2106.02374>
- Jonathan Dumas, Bertrand Cornélusse, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Stochastic and deterministic formulations for capacity firming nominations. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–7. IEEE, 2020. URL <https://arxiv.org/abs/2106.02425>
- Jonathan Dumas, Colin Cointe, Xavier Fettweis, and Bertrand Cornélusse. Deep learning-based multi-output quantile forecasting of pv generation. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494976. URL <https://arxiv.org/abs/2106.01271>
- Jonathan Dumas, Bertrand Cornélusse, Xavier Fettweis, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Probabilistic forecasting for sizing in the capacity firming framework. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494947. URL <https://arxiv.org/abs/2106.02323>
- Jonathan Dumas, Colin Cointe, Antoine Wehenkel, Antonio Sutera, Xavier Fettweis, and Bertrand Cornélusse. A probabilistic forecast-driven strategy for a risk-aware participation in the capacity firming market. *IEEE Transactions on Sustainable Energy*, pages 1–1, 2021. doi: 10.1109/TSTE.2021.3117594. URL <https://arxiv.org/abs/2105.13801>

-
- Jonathan Dumas, Antoine Wehenkel, Damien Lanaspeze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

Part I

Forecasting

Overview

Part I presents the forecasting techniques and metrics required to produce and evaluate reliable point and probabilistic forecasts to be used as input of decision-making models in Part II. Then, it investigates the various types of forecasts in several case studies: point forecasts, quantile forecasts, prediction intervals, density forecasts, and scenarios.

“ I never think of the future — it comes soon enough.”

— **Albert Einstein**

“ We have two classes of forecasters: Those who don't know — and those who don't know they don't know.”

— **John Kenneth Galbraith**

Figure 1.3 illustrates the organization of Part I, which can be read in two passes depending on the forecasting knowledge. First, a forecasting practitioner may identify the forecasting type of interest and select the corresponding Chapter. For instance, Chapter 6 studies the scenarios of renewable generation and electrical consumption. Second, a forecasting entrant should be interested in reading Chapters 2 and 3 to acquire the forecasting basics. Then, the following Chapters are the application of each type of forecast on case studies.

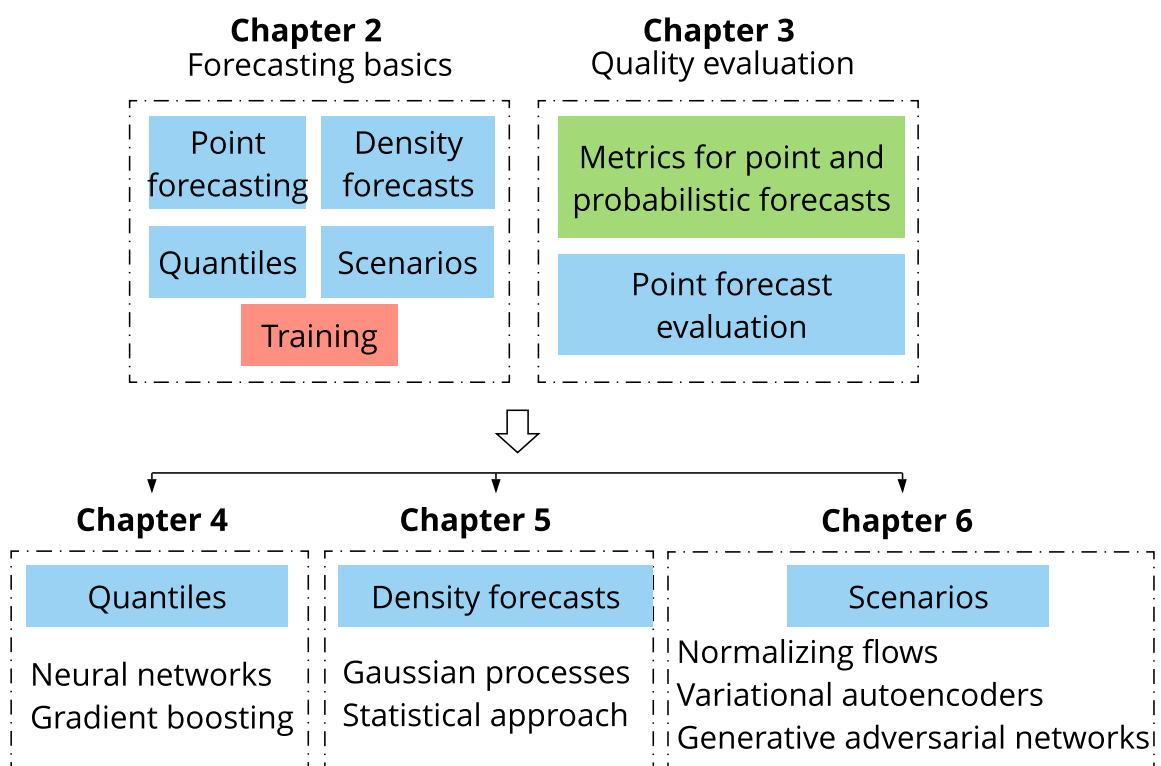


Fig. 1.3 Part I skeleton.

Part I general nomenclature

Acronyms

Name	Description
PDF	Probability Density Function.
CDF	Cumulative Distribution Function.
(N)MAE	(Normalized) mean absolute error.
(N)RMSE	(Normalized) root mean square error.
CRPS	Continuous ranked probability score.
IS	Interval score.
ES	Energy score.
VS	Variogram score.
DM	Diebold-Mariano.
RNN	Recurrent Neural Network.
LSTM	Long Short-Term Memory.
BLSTM	Bidirectional LSTM.
GBR	Gradient boosting regression.
MLP	Multi-layer perceptron.
ED	Encoder-decoder.
NFs	Normalizing Flows.
GANs	Generative adversarial Networks.
VAEs	Variational AutoEncoders.

Variables and parameters

Name	Range	Description
X		Continuous random variable.
\mathbf{X}		Continuous multivariate random variable.
x	$\in \mathbb{R}$	Realization of the random variable X .

\mathbf{x}	$\in \mathbb{R}^T$	Realization of the multivariate random variable \mathbf{X} .
\hat{x}	$\in \mathbb{R}$	Point forecast of x .
$\hat{\mathbf{x}}$	$\in \mathbb{R}^T$	Multi-output point forecast of \mathbf{x} .
$\hat{x}^{(q)}$	$\in \mathbb{R}$	Quantile forecast of x .
$\hat{\mathbf{x}}^{(q)}$	$\in \mathbb{R}^T$	Multi-output quantile forecast of \mathbf{x} .
$\hat{I}^{(\alpha)}$	$\in \mathbb{R}^2$	Prediction interval with a coverage rate of $(1 - \alpha)$.
$\hat{\mathbf{I}}^{(\alpha)}$	$\in \mathbb{R}^{2T}$	Multi-output prediction interval with a coverage rate of $(1 - \alpha)$.
\hat{x}^i	$\in \mathbb{R}$	Scenario i of x .
$\hat{\mathbf{x}}^i$	$\in \mathbb{R}^T$	Scenario i of \mathbf{x} .
ϵ	$\in \mathbb{R}$	Point forecast error.
ξ	$\in \{0, 1\}$	Indicator variable.
$\tilde{\xi}$	$\in \mathbb{R}^+$	Empirical level.

Symbols

Name	Description
g_θ	Forecasting model with parameters θ .
\mathbb{E}	Expectation.
f	Probability Density Function.
F	Cumulative Distribution Function.
\hat{f}	Density forecast of the pdf.
\hat{F}	Density forecast of the cdf.
ρ	Pinball loss function.

Sets and indices

Name	Description
t	Time index.
k	Forecasting lead time index.
T	Forecasting horizon.
q	Quantile index.
ω	Scenario index.
$\#\Omega$	Number of scenarios.
\mathcal{T}	Set of time periods, $\mathcal{T} = \{1, 2, \dots, T\}$.
Ω	Set of scenarios, $\Omega = \{1, 2, \dots, \#\Omega\}$.
\mathcal{D}	Information set.

Chapter 2

Forecasting background

Overview

This chapter introduces the main concepts in forecasting as a background for the work developed in the following chapters of this manuscript. It presents the various types of forecasts: point forecasts, quantile forecasts, prediction intervals, density forecasts, and scenarios. In addition, it provides some knowledge on how to train a forecasting model.

General textbooks [124, 83, 184, 75] provide further information for the interested reader. Two courses on this topic also provide interesting material: (1) "Renewables in Electricity Markets"^a given by professor Pierre Pinson at the Technical University of Denmark. It covers some basics of electricity markets, the impact of renewables on markets, participation of renewable energy producers in electricity markets, and renewable energy analytics (mainly forecasting); (2) "INFO8010 - Deep Learning"^b, ULiège, Spring 2021, given by associate professor Gilles Louppe at Liège University. It covers the foundations and the landscape of deep learning.

^a<http://pierrepinson.com/index.php/teaching/>

^b<https://github.com/glouppe/info8010-deep-learning>

“*If life were predictable it would cease to be life, and be without flavor.*”

— Eleanor Roosevelt

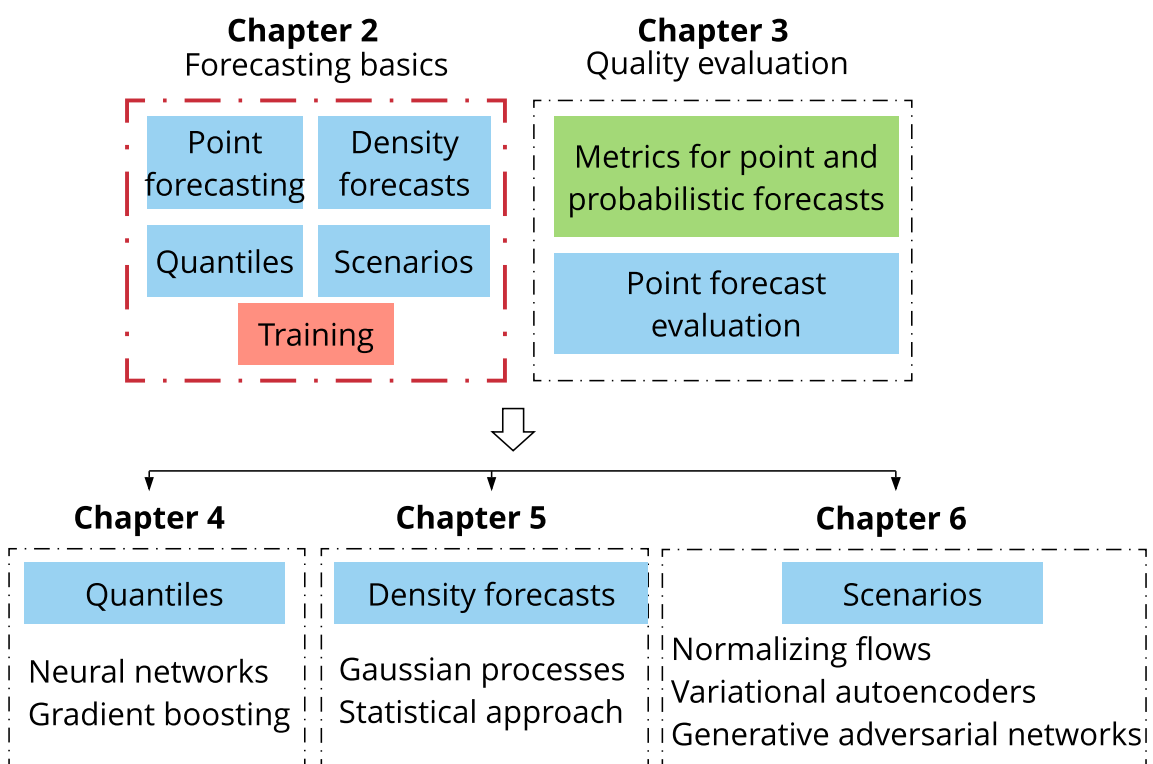


Fig. 2.1 Chapter 2 position in Part I.

Following Morales et al. [124] power generation from renewable energy sources, such as wind and solar, are referred to as *stochastic power generation* in this thesis. Electrical consumption and electricity prices are also modeled as stochastic variables. Predictions of renewable energy generation, consumption, and electricity prices can be obtained and presented differently. The choice of a particular kind of forecast depends on the process characteristics of interest to the decision-maker and the type of operational problem. The various types of forecasts and their presentation are introduced in the following, starting from the most common *point forecasts* and building up towards the more advanced products that are *probabilistic forecasts* and scenarios.

This Chapter is organized as follows. Section 2.1 introduces the point forecasts. Section 2.2 presents the various types of probabilistic forecasts. Section 2.3 proposes an abstract formulation of a model-based forecaster. Section 2.4 provides some knowledge on how to train a forecasting model. Finally, conclusions are drawn in Section 2.5.

2.1 Point forecast

Let $x_t \in \mathbb{R}$ be the variable of interest, *e.g.*, renewable energy generation, consumption, or electricity prices, measured at time t , which corresponds to a realization of the random variable X_t .

Definition 2.1.1 (A model-based forecast). [124, Chapter 2] A (model-based) forecast $\hat{x}_{t+k|t} \in \mathbb{R}$ is an estimate of some of the characteristics of the stochastic process X_{t+k} issued at time t for time $t+k$ given a model g_θ , with parameters θ and the information set \mathcal{D} gathering all data and knowledge about the processes of interest up to time t , such as weather forecasts, historical observations, calendars variables, *etc.*

In the above definition, k is the *lead time*, sometimes also referred to as *forecast horizon*. The 'hat' symbol expresses that $\hat{x}_{t+k|t}$ is an estimate only. It models the presence of uncertainty both in our knowledge of the process and inherent to the process itself. The forecast for time $t+k$ is conditional on our knowledge of stochastic process up to time t , including the data used as input to the forecasting process and the models identified and parameters estimated. Therefore, a forecaster somewhat makes the crucial assumption that the future will be like the past.

Forecasts are series of consecutive values $\hat{x}_{t+k|t}$, $k = k_1, \dots, k_T$, that is, for regularly spaced lead times up to the *forecast length* T . That regular spacing Δt is called the *temporal resolution* of the forecasts. For instance, when one talks of day-ahead forecasts with an hourly resolution, forecasts consist of a series gathering predicted power values for each of the following 24 hours of the next day.

Definition 2.1.2 (Point forecast). [124, Chapter 2] A *point forecast* $\hat{x}_{t+k|t} \in \mathbb{R}$ is a single-valued issued at time t for $t+k$, and corresponds to the conditional expectation of X_{t+k}

$$\hat{x}_{t+k|t} := \mathbb{E}[X_{t+k|t}|g_\theta, \mathcal{D}], \quad (2.1)$$

given g_θ , and the information set \mathcal{D} .

A forecast in the form of a conditional expectation translates to acknowledging the presence of uncertainty, even though it is not quantified and communicated.

Definition 2.1.3 (Multi-output point forecast). A *multi-output point forecast* computed at t for $t+k_1$ to $t+k_T$ is the vector

$$\hat{\mathbf{x}}_t := [\hat{x}_{t+k_1|t}, \dots, \hat{x}_{t+k_T|t}]^\top \in \mathbb{R}^T. \quad (2.2)$$

Depending on the problem formulation, it can be computed directly as a vector or as an aggregate of single output forecasts.

2.2 Probabilistic forecasts

In contrast to point predictions, probabilistic forecasts aim at providing decision-makers with the full information about potential future outcomes. Let f_t and F_t be the *probability density function* (PDF) and related *cumulative distribution function* (CDF) of X_t , respectively.

Definition 2.2.1 (Probabilistic forecast). [124, Chapter 2] A probabilistic forecast issued at time t for time $t+k$ consists of a prediction of the PDF (or equivalently, the CDF) of X_{t+k} , or of some summary features.

Various types of probabilistic forecasts have been developed: *quantile*, *prediction intervals*, *scenarios*, and *density* forecasts.

2.2.1 Quantiles

Definition 2.2.2 (Quantile forecast). [124, Chapter 2] A *quantile forecast* $\hat{x}_{t+k|t}^{(q)} \in \mathbb{R}$ with nominal level q is an estimate, issued at time t for time $t+k$ of the quantile $x_{t+k|t}^{(q)}$ for the random variable $X_{t+k|t}$

$$P[X_{t+k|t} \leq \hat{x}_{t+k|t}^{(q)} | g_\theta, \mathcal{D}] = q, \quad (2.3)$$

given g_θ , and the information set \mathcal{D} . Or equivalently $\hat{x}_{t+k|t}^{(q)} = \hat{F}_{t+k|t}^{-1}(q)$, with \hat{F} the estimated cumulative distribution function of the continuous random variable X .

By issuing a quantile forecast $\hat{x}_{t+k|t}^{(q)}$, the forecaster tells at time t that there is a probability q that x_{t+k} will be less than $\hat{x}_{t+k|t}^{(q)}$ at time $t+k$. Quantile forecasts are of interest for several operational problems. For instance, the optimal day-ahead bidding of wind or PV generation uses quantile forecasts whose nominal level is a simple function of day-ahead and balancing market prices [23]. Furthermore, quantile forecasts also define prediction intervals that can be used for robust optimization.

Definition 2.2.3 (Multi-output quantile forecast). A *multi-output quantile forecast* of length T with nominal level q computed at t for $t+k_1$ to $t+k_T$ is the vector

$$\hat{\mathbf{x}}_t^{(q)} := [\hat{x}_{t+k_1|t}^{(q)}, \dots, \hat{x}_{t+k_T|t}^{(q)}]^\top \in \mathbb{R}^T. \quad (2.4)$$

2.2.2 Prediction intervals

Quantile forecasts give probabilistic information in the form of a threshold level associated with a probability. Even though they may be of direct use for several operational problems, they cannot provide forecast users with a feeling about the level of forecast uncertainty for the coming period. For that purpose, prediction intervals define the range of values within which the observation is expected to be with a certain probability, *i.e.*, its nominal coverage rate [144].

Definition 2.2.4 (Prediction interval). [124, Chapter 2] A *prediction interval* $\hat{I}_{t+k|t}^{(\alpha)} \in \mathbb{R}^2$ issued at t for $t+k$, defines a range of potential values for X_{t+k} , for a certain level of probability $(1-\alpha)$, $\alpha \in [0, 1]$. Its nominal coverage rate is

$$P[X_{t+k} \in \hat{I}_{t+k|t}^{(\alpha)} | g_\theta, \mathcal{D}] = 1 - \alpha. \quad (2.5)$$

Definition 2.2.5 (Central prediction interval). [124, Chapter 2] A *central prediction interval* consists of centering the prediction interval on the median where there is the same probability of risk below and above the median. A central prediction interval with a coverage rate of $(1-\alpha)$ is estimated by using the quantiles $q = (\alpha/2)$ and $q = (1-\alpha/2)$. Its nominal coverage rate is

$$\hat{I}_{t+k|t}^{(\alpha)} = [\hat{x}_{t+k|t}^{(q=\alpha/2)}, \hat{x}_{t+k|t}^{(q=1-\alpha/2)}]. \quad (2.6)$$

For instance, central prediction interval with a nominal coverage rate of 90%, *i.e.*, $(1-\alpha) = 0.9$, are defined by quantile forecasts with nominal levels of 5 and 95%.

Definition 2.2.6 (Multi-output central prediction interval). A *multi-output central prediction interval* with a coverage rate of $(1 - \alpha)$ computed at t for $t + k_1$ to $t + k_T$ is the vector

$$\hat{\mathbf{I}}_{t+k|t}^{(\alpha)} = [\hat{I}_{t+k_1|t}^{(\alpha)}, \dots, \hat{I}_{t+k_T|t}^{(\alpha)}]^\top \in \mathbb{R}^{2T}. \quad (2.7)$$

2.2.3 Scenarios

Let us introduce the *multivariate* random variable

$$\mathbf{X}_t := \{X_{t+k}\}, \quad k = k_1, \dots, k_T, \quad (2.8)$$

which gathers the random variables characterizing the stochastic power generation process for the T following lead times. Hence, it covers their marginal densities as well as their interdependence structure.

Definition 2.2.7 (Scenarios). *Scenarios* issued at time t and for a set of T successive lead times, *i.e.*, with $k = k_1, \dots, k_T$ consist of a set of M time trajectories

$$\hat{\mathbf{x}}_t^i := [\hat{x}_{t+k_1|t}^i, \dots, \hat{x}_{t+k_T|t}^i]^\top \in \mathbb{R}^T \quad i = 1, \dots, M. \quad (2.9)$$

The resulting time trajectories comprise scenarios like those commonly used in stochastic programming.

2.2.4 Density forecasts

All the various types of predictions presented in the above, *i.e.*, point, quantile, and interval forecasts, are only partly describing the complete information about the future of X at every lead time. Density forecasts would give this whole information for each point of time in the future.

Definition 2.2.8 (Density forecast). [124, Chapter 2] A *density forecasts* $\hat{f}_{t+k|t}$ ($\hat{F}_{t+k|t}$) issued at time t for $t+k$, is a complete description of the pdf (or cdf) of X_{t+k} conditional on a given model g_θ , and the information set \mathcal{D} .

2.3 Model-based formulation

Let assume the information set $\mathcal{D} := \{\mathbf{x}_t, \mathbf{c}_t\}_{t=1}^N$ is composed of N independent and identically distributed samples from the joint distribution $p(\mathbf{x}, \mathbf{c})$ of two continuous variables X and C . X is the variable of interest, *e.g.*, renewable energy generation,

consumption, or electricity prices, and C is the context, *e.g.*, the weather forecasts, calendar variables, or exogenous variables. Generically, any prediction of a random variable X issued at time t , being point or probabilistic forecast is a linear or nonlinear function of \mathcal{D} . The goal of Part I is to generate multi-output context-based forecasts $\hat{\mathbf{x}}$ that are distributed under $p(\mathbf{x}|\mathbf{c})$.

Definition 2.3.1 (Multi-output model-based forecasts). A *multi-output model-based forecasts* of length T computed by g_θ at t for $t + k_1$ to $t + k_T$ is the vector

$$\hat{\mathbf{x}}_t \sim g_\theta(\mathbf{x}_{<t}, \mathbf{c}_t) \in \mathbb{R}^T, \quad (2.10)$$

given the context \mathbf{c}_t , and observations $\mathbf{x}_{<t} = [x_1, \dots, x_{t-1}]^\top$ up to time t .

Its purpose is to generate synthetic but realistic data $\hat{\mathbf{x}}_t$ whose distribution is as close as possible to the unknown data distribution $p(\mathbf{x}|\mathbf{c})$. When considering point forecasts, quantile forecasts, and scenarios, $\hat{\mathbf{x}}_t$ is defined by (2.2), (2.4), and (2.9), respectively. This abstract formulation is used in Section 3.3 and Chapters 4 and 6 where the forecasting models and the related inputs including the context are specified.

2.4 Model training

This section provides the basics of supervised learning that is used in Part I to train the forecasting models g_θ . It relies mainly on Lecture 1 of INFO8010 - Deep Learning¹, ULiège, Spring 2021, from associate professor Gilles Louppe [119] at Liège University. The interested reader may find interesting materials in Duchesne [47, Chapter 3] and Sutera [162, Chapter 2]. They introduce the different types of machine learning problems, describe their characteristics, and the procedure to apply supervised learning.

2.4.1 Regression with supervised learning

Consider the unknown joint probability distribution $p(\mathbf{x}, \mathbf{c})$ of two continuous variables X , *e.g.*, renewable energy generation, and C , *e.g.*, the weather forecasts, introduced in the previous Section. Let assume some training data $\mathcal{D} = \{\mathbf{x}_t, \mathbf{c}_t\}_{t=1}^N$ composed of N independent and identically distributed samples.

Supervised learning is usually concerned with the two following *inference* problems: *classification* and *regression*. Classification consists of identifying a decision boundary

¹<https://github.com/glouppe/info8010-deep-learning>

between objects of distinct classes. Regression aims at estimating relationships among (usually continuous) variables. In this thesis, we focus on regression.

Definition 2.4.1 (Regression). Given $\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}$, for $t = 1, \dots, N$ we would like to estimate for any new \mathbf{c}

$$\mathbb{E}[X = \mathbf{x} | C = \mathbf{c}]. \quad (2.11)$$

For instance, let assume g_θ is parameterized with a neural network where the last layer does not contain any final activation. If we make the assumption that $p(\mathbf{x}, \mathbf{c}) \sim \mathcal{N}(\mathbf{x}; \mu = g_\theta(\mathbf{x}|\mathbf{c}); \sigma^2 = 1)$, we can perform maximum likelihood estimation to estimate the model's parameters θ that is equivalent to minimize the negative log likelihood

$$\arg \max_{\theta} p(\mathcal{D}|\theta) \sim \arg \min_{\theta} - \log \prod_{\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}} p(\mathbf{x} = \mathbf{x}_t | \mathbf{c}_t, \theta) \quad (2.12a)$$

$$= \arg \min_{\theta} - \sum_{\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}} \log p(\mathbf{x} = \mathbf{x}_t | \mathbf{c}_t, \theta) \quad (2.12b)$$

$$= \arg \min_{\theta} - \sum_{\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}} \log \left(\frac{1}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} \|\mathbf{x}_t - g_\theta(\mathbf{x}_t | \mathbf{c}_t)\|^2 \right) \right) \quad (2.12c)$$

$$\sim \arg \min_{\theta} \sum_{\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}} \|\mathbf{x}_t - g_\theta(\mathbf{x}_t | \mathbf{c}_t)\|^2, \quad (2.12d)$$

which recovers the common squared error loss $L(\mathbf{x}, \hat{\mathbf{x}} \sim g_\theta(\mathbf{x}|\mathbf{c})) = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$ that will be used in Section 3.3 for point forecasting.

2.4.2 Empirical risk minimization

Consider a function g_θ produced by a learning algorithm. The predictions of this function can be evaluated through a loss $L : \mathbb{R}^T \times \mathbb{R}^T \rightarrow \mathbb{R}$, such that $L(\mathbf{x}, \hat{\mathbf{x}} \sim g_\theta(\mathbf{x}|\mathbf{c})) \geq 0$ measures how close the prediction $\hat{\mathbf{x}}$ from \mathbf{x} is. For instance, in point forecasting L is the mean squared error or the pinball loss for quantile forecasting.

The key idea of the model training relies on the *empirical risk minimization*. Let \mathcal{G} denote the hypothesis space, *i.e.* the set of all functions g_θ than can be produced by the chosen learning algorithm.

Definition 2.4.2 (Empirical risk minimization). [169] We are looking for a function g_θ with a small expected risk

$$R(g_\theta) = \mathbb{E}_{\mathbf{x}, \mathbf{c} \sim p(\mathbf{x}, \mathbf{c})} [L(\mathbf{x}, \hat{\mathbf{x}} \sim g_\theta(\mathbf{x}|\mathbf{c}))], \quad (2.13)$$

also called the generalization error.

Therefore, for a given data generating distribution $p(\mathbf{x}, \mathbf{c})$ and for a given hypothesis space \mathcal{G} , the optimal model is

$$g_\theta^* = \arg \min_{g_\theta \in \mathcal{G}} R(g_\theta). \quad (2.14)$$

However, since $p(\mathbf{x}, \mathbf{c})$ is unknown, the expected risk cannot be evaluated and the optimal model cannot be determined. Nevertheless, if we have some training data $\mathcal{D} = \{\mathbf{x}_t, \mathbf{c}_t\}_{t=1}^N$ composed of N independent and identically distributed samples, we can compute an estimate that is the empirical risk [169] (or training error)

$$\hat{R}(g_\theta, \mathcal{D}) = \frac{1}{N} \sum_{\mathbf{x}_t, \mathbf{c}_t \in \mathcal{D}} L(\mathbf{x}_t, \hat{\mathbf{x}}_t \sim g_\theta(\mathbf{x}_t|\mathbf{c}_t)). \quad (2.15)$$

This estimator is unbiased and can be used for finding a good enough approximation of g_θ^* , resulting in the empirical risk minimization principle

$$g_{\theta, \mathcal{D}}^* = \arg \min_{g_\theta \in \mathcal{G}} \hat{R}(g_\theta, \mathcal{D}). \quad (2.16)$$

Note: most machine learning algorithms, including neural networks, implement empirical risk minimization. Under regularity assumptions, empirical risk minimizers converge: $\lim_{N \rightarrow \infty} g_{\theta, \mathcal{D}}^* \rightarrow g_\theta^*$.

The capacity of a hypothesis space induced by a learning algorithm intuitively represents the ability to find a suitable model $g_\theta \in \mathcal{G}$ for any function, regardless of its complexity. In practice, capacity can be controlled through hyper-parameters θ of the learning algorithm. Then, the goal is to adjust the capacity of the hypothesis space \mathcal{G} such that the expected risk of the empirical risk minimizer gets as low as possible. To this end, it is essential to understand the concept of the bias-variance trade-off [64]. First, reducing the capacity makes $g_{\theta, \mathcal{D}}^*$ fit the data less on average, which increases the bias term (under-fitting). Second, increasing the capacity makes $g_{\theta, \mathcal{D}}^*$ vary a lot with the training data, which increases the variance term (over-fitting). Therefore, the bias-variance trade-off implies that a model should balance under-fitting

and over-fitting: rich enough to express underlying structure in data, simple enough to avoid fitting spurious patterns. It is summarized in the classical U-shaped risk curve [9], shown in Figure 2.2.




Fig. 2.2 The classical U-shaped risk curve arising from the bias-variance trade-off. Curves for training risk (dashed line) and test risk (solid line). Note: \mathcal{H} on the Figure is \mathcal{G} . Credits: [9].

2.4.3 The "double descent" curve

In the past few years, several studies have raised questions about the mathematical foundations of machine learning and their relevance to practitioners. Indeed, very complex machine learning models such as neural networks are trained to exactly fit the data. Classically, following the concept of bias-variance trade-off, these models should be considered over-fit. However, they often obtain high accuracy on test data. It is illustrated by Figure 2.3 from Belkin et al. [9] that discuss empirical evidence for the double descent curve. This study demonstrated that the capacity of the function

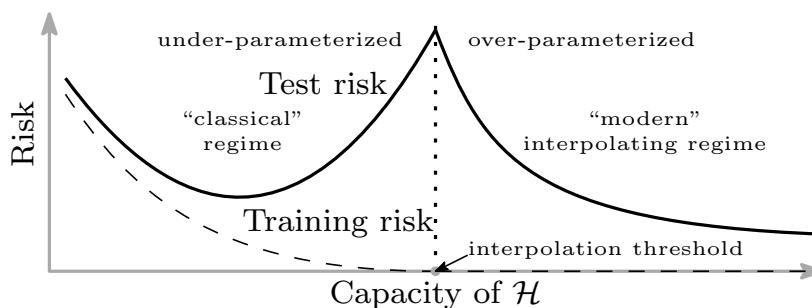


Fig. 2.3 The double descent risk curve, which incorporates the U-shaped risk curve (*i.e.*, the "classical" regime) together with the observed behavior from using high capacity function classes (*i.e.*, the "modern" interpolating regime), separated by the interpolation threshold. Credits: [9].

class does not necessarily reflect how well the predictor matches the *inductive bias*

appropriate for the problem at hand. By training models on a range of real-world datasets and synthetic data, the inductive bias that seems appropriate is the regularity or smoothness of a function as measured by a specific function space norm. Then, by considering larger function classes, which include more candidate predictors compatible with the data, the training allowed to determine interpolating functions with smaller norms and are thus "simpler". Therefore, increasing function class capacity improves the performance of classifiers. This connection investigated by Belkin et al. [9] between the performance and the structure of machine learning models delineates the limits of classical analyses. It has implications for both the theory and practice of machine learning.

2.4.4 Training methodology

In practice, one has a finite dataset of input-output pairs \mathcal{D} and no further information about the joint probability distribution $p(\mathbf{x}, \mathbf{c})$. Then, the model is selected by minimizing the empirical risk (2.15) over the dataset. Therefore, it is typically strongly biased in an optimistic way and is a bad estimate of the generalization error, also called the testing error. A "good" model should predict well data independent from the dataset used for training but drawn from the same distribution.

2.4.5 Learning, validation, and testing sets

A *good practice* in machine learning is to use a dedicated part of the dataset as an independent testing set that is not used to train the models but only to estimate the generalization error. Ideally, the dataset \mathcal{D} is divided randomly into three parts, as depicted in Figure 2.4.

1. The models are trained on the learning set.
2. The validation set is used: (1) to evaluate the generalization error of the trained models to select among several learning algorithms the one more suited to the studied problem; (2) to optimize some algorithm's hyper-parameters θ and to avoid over-fitting.
3. Finally, the testing set is kept until the end of the process. It allows assessing the performance of the selected model on independent data.

A common rule is to build the training set as large as possible to obtain good predictors while keeping enough samples in the validation and testing sets to correctly conduct the hyper-parameters selection and estimate the generalization error properly.




Fig. 2.4 Proper evaluation protocols. Credits: Francois Fleuret, EE559 Deep Learning, EPFL <https://fleuret.org/dlc/>.

2.4.6 k -fold cross-validation

When the dataset is composed of only a few months of data, dividing it into three parts can lead to a too "small" learning set to learn good predictors regarding both the empirical risk and the generalization error. In this case, the k -fold cross-validation methodology allows evaluating the generalization error of the algorithm correctly [83].

It consists of dividing the dataset into k -folds. Then, the model is trained on $k - 1$ folds, and one fold is left out to evaluate the testing error. In total, the model is trained k times with k pairs of learning and testing sets. The generalization error is estimated by averaging the k errors computed on k testing sets. This procedure is adopted in Chapters 4 and 5 where the dataset is composed of only a few months. Figure 2.5 depicts a 5-folds cross-validation.

However, k -fold cross-validation may suffer from high variability, which can be responsible for bad choices in model selection and erratic behavior in the estimated expected prediction error. A study conducted by Bengio and Grandvalet [13] demonstrated there is no unbiased estimator of the variance of k -fold cross-validation. Therefore, the assessment of the significance of observed differences in cross-validation scores should be treated with caution.

2.5 Conclusions

This Chapter introduces the forecasting basics with the various types of forecasts, starting from the most common point forecasts and building up towards the more advanced probabilistic products: quantiles, prediction intervals, scenarios, and density forecasts. It also provides some basics on how to train and evaluate a forecasting model properly. The next Chapter presents the methodologies to evaluate the quality of these various types of forecasts.

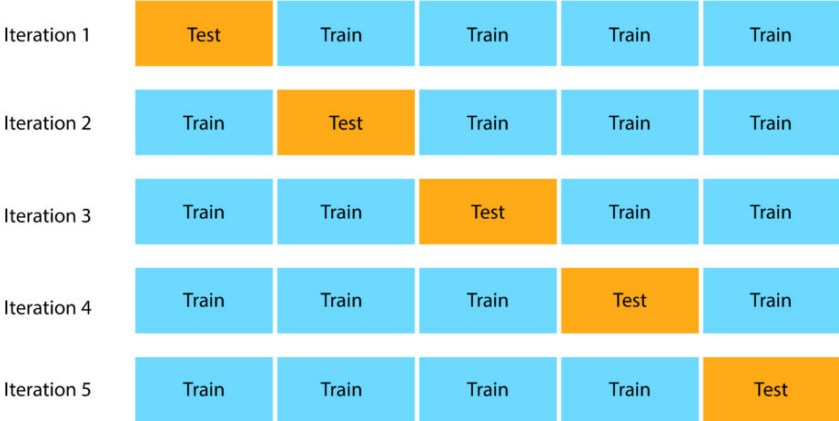


Fig. 2.5 *k*-fold cross-validation.

Chapter 3

Forecast evaluation

Overview

This chapter introduces the main concepts in forecasting evaluation as a background for the work developed in the following chapters of this manuscript.

General textbooks such as [124] provide further information to the interested reader. In addition, the course "Renewables in Electricity Markets"^a given by professor Pierre Pinson at the Technical University of Denmark proposes valuable materials on this topic.

^a<http://pierrepinson.com/index.php/teaching/>

“ Everything that can be counted does not necessarily count; everything that counts cannot necessarily be counted. ”

— **Albert Einstein**

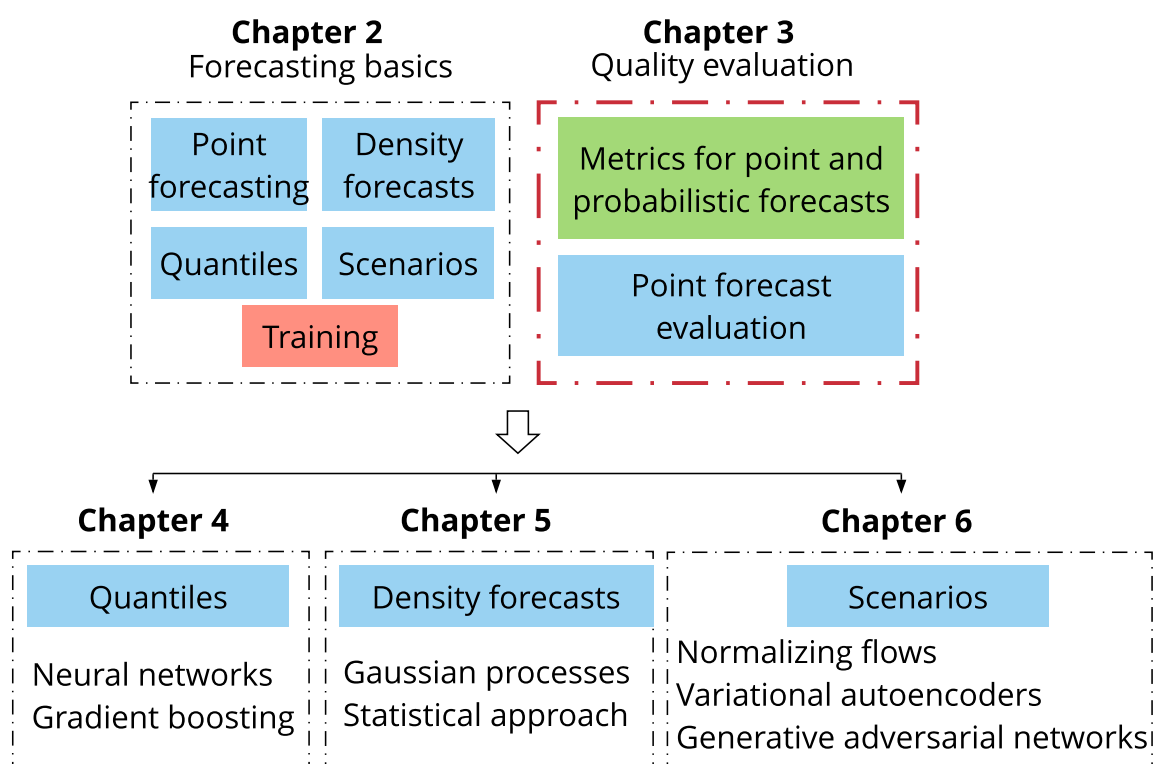


Fig. 3.1 Chapter 3 position in Part I.

For predictions in any form, one must differentiate between their *quality* and their *value* [124]. The forecast quality corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. The forecast value, studied in Part II, relates to the benefits of using forecasts in decision-making, such as participation in the electricity market. Consequently, forecast quality is independent of the current operational problem while not forecast value. Intuitively, it is relevant to evaluate the forecast quality before using the predictions as input to operational problems. Even if a good quality does not necessarily imply a good forecast value, the quality assessment provides interesting information about the model at hand.

This Chapter adopts the general framework of Morales et al. [124] to provide the basics required to appraise the quality of predictions. It is organized as follows. Sections 3.1 and 3.2 introduce the metrics for point and probabilistic forecasting, respectively. Finally, conclusions are drawn in Section 3.4.

3.1 Metrics for point forecasts

The base quantity for evaluating point forecasts of a continuous random variable is the forecast error.

Definition 3.1.1 (Forecast error). [124, Chapter 2] The *forecast error* $\epsilon_{t+k|t}$ is the difference between observed and predicted values

$$\epsilon_{t+k|t} := x_{t+k|t} - \hat{x}_{t+k|t}. \quad (3.1)$$

Note that the forecast error may be normalized so that verification results can be comparable for different time series. If so, normalization is most commonly performed by the nominal capacity of the site of interest. The first error criterion that may be computed is the bias of the forecasts, which corresponds to the systematic part of the forecast error. It may be corrected in a straightforward manner using simple statistical models.

Definition 3.1.2 (Bias). [124, Chapter 2] The *bias* is the mean of all errors over the evaluation period of length N , considered indifferently. For lead time k , it is

$$\text{bias}(k) := \frac{1}{N} \sum_{t=1}^N \epsilon_{t+k|t}. \quad (3.2)$$

This summary measure does not tell much about the quality of point forecasts but only about a systematic error that should be corrected. Therefore, for a better appraisal of the forecasts, it is advised to use scores such as the mean absolute and root mean square errors.

Definition 3.1.3 (Mean absolute error). [124, Chapter 2] The *mean absolute error* (MAE) is defined as the average of the absolute forecast errors over an evaluation set of length N

$$\text{MAE}(k) := \frac{1}{N} \sum_{t=1}^N |\epsilon_{t+k|t}|. \quad (3.3)$$

Definition 3.1.4 (Root mean square error). [124, Chapter 2] The *root mean square error* (RMSE) is defined as the square root of the sum of squared errors over an evaluation of length N

$$\text{RMSE}(k) := \sqrt{\frac{1}{N} \sum_{t=1}^N (\epsilon_{t+k|t})^2}. \quad (3.4)$$

All the above error criteria are independent of the length of the evaluation set. The nominal capacity of the renewable energy site of interest can be used to normalize them, and then they are referred to Nbias, NRMSE, and NMAE.

3.2 Metrics for probabilistic forecasts

3.2.1 Calibration

The first requirement of probabilistic forecasts is to consistently inform about the probability of events. It leads to the concept of *probabilistic calibration*, also referred to as *reliability*. The assessment of probabilistic calibration only informs about a form of bias of probabilistic forecasts. A frequentist approach, based on an evaluation set of sufficient length, can be performed by assessing the reliability of each of the defining quantile forecasts using the indicator variable.

Definition 3.2.1 (Indicator variable). [124, Chapter 2] The *indicator variable* $\xi_{t,k}^q$, for a given quantile forecast $\hat{x}_{t+k|t}^{(q)}$ and corresponding realization x_{t+k} is

$$\xi_{t,k}^q := 1_{\{x_{t+k} < \hat{x}_{t+k|t}^{(q)}\}} = \begin{cases} 1 & x_{t+k} < \hat{x}_{t+k|t}^{(q)} \\ 0 & x_{t+k} \geq \hat{x}_{t+k|t}^{(q)} \end{cases}. \quad (3.5)$$

$\xi_{t,k}^q$ is a binary variable indicating if the quantile forecasts cover, or not, the measurements. The empirical level of quantile forecasts can be defined, estimated, and eventually compared with their nominal one by using this indicator variable.

Definition 3.2.2 (Empirical level). [124, Chapter 2] The *empirical level* $\tilde{\xi}_{t,k}^q$, for a nominal level q and lead time k , is obtained by calculating the mean of the $\{\xi_{t,k}^q\}_{t=1}^N$ time-series over an evaluation set of length N

$$\tilde{\xi}_{t,k}^q := \frac{1}{N} \sum_{t=1}^N \xi_{t,k}^q. \quad (3.6)$$

The difference between nominal and empirical levels of quantile forecasts is to be seen as a probabilistic bias. Then, probabilistic calibration may be appraised visually by using *reliability diagrams* plotting empirical *vs.* nominal levels of the quantiles defining density forecasts.

3.2.2 Univariate skill scores

Perfectly calibrated probabilistic forecasts do not guarantee that the forecasts are "good". For instance, they may not discriminate among situations with various uncertainty levels, while these aspects are crucial in decision-making. The overall quality of probabilistic forecasts may be assessed based on skill scores. First, we consider the *univariate* skill score that can only assess the quality of the forecasts with respect to their marginals (time period). Second, we focus on the *multivariate* skill score that can directly assess multivariate scenarios.

Continuous ranked probability score

The continuous ranked probability score (CRPS) [70] is a univariate scoring rule that penalizes the lack of resolution of the predictive distributions as well as biased forecasts. It is negatively oriented, *i.e.*, the lower, the better, and for deterministic forecasts, it turns out to be the mean absolute error (MAE). Thus, it can be directly compared to the MAE criterion used for point forecasts since the CRPS is its generalization in a probabilistic forecasting framework. The CRPS is used to compare the skill of predictive marginals for each component of the random variable of interest. In our case, for the twenty-four time periods of the day.

Definition 3.2.3 (CRPS-integral). [70] The *CRPS* for predictive densities $\hat{F}_{t+k|t}$ and corresponding measurement $x_{t+k|t}$, is calculated over an evaluation set of length N

$$\text{CRPS}(k) := \frac{1}{N} \sum_{t=1}^N \int_{x'} [\hat{F}_{t+k|t}(x') - 1_{\{x' > x_{t+k|t}\}}]^2 dx'. \quad (3.7)$$

However, it is not easy to estimate the integral form of the CRPS defined by (3.7).

Definition 3.2.4 (CRPS-energy). [70] Gneiting and Raftery [70] propose a formulation called the *energy form* of the CRPS since it is just the one-dimensional case of the energy score, defined in negative orientation as follows

$$\text{CRPS}(P, x_{t+k|t}) = \mathbb{E}_P[|X - x_{t+k|t}|] - \frac{1}{2} \mathbb{E}_P[|X - X'|], \quad (3.8)$$

where X and X' are independent random variables with distribution P and finite first moment, and \mathbb{E}_P is the expectation according to the probabilistic distribution P .

The CRPS can be computed for quantile forecasts and scenarios. Let $\{\hat{x}_{t+k|t}^i\}_{i=1}^M$ be the set of M scenarios generated at time t for lead time k . The estimator of (3.8) proposed by Zamo and Naveau [181], over an evaluation set of length N for lead time k , is

$$\text{CRPS}(k) = \frac{1}{N} \left[\sum_{t=1}^N \frac{1}{M} \sum_{i=1}^M |\hat{x}_{t+k|t}^i - x_{t+k|t}| - \frac{1}{2M^2} \sum_{i,j=1}^M |\hat{x}_{t+k|t}^i - \hat{x}_{t+k|t}^j| \right]. \quad (3.9)$$

Let $\{\hat{x}_{t+k|t}^{(q)}\}_{q=1}^Q$ be the set of Q quantiles generated at time t for lead time k . (3.8) is over an evaluation set of length N for lead time k

$$\text{CRPS}(k) = \frac{1}{N} \left[\sum_{t=1}^N \frac{1}{Q} \sum_{q=1}^Q |\hat{x}_{t+k|t}^{(q)} - x_{t+k|t}| - \frac{1}{2Q^2} \sum_{q,q'=1}^Q |\hat{x}_{t+k|t}^{(q)} - \hat{x}_{t+k|t}^{(q')}| \right]. \quad (3.10)$$

Quantile score

The quantile score (QS), also known as the *pinball loss*, is complementary to the CRPS. It permits obtaining detailed information about the forecast quality at specific probability levels, *i.e.*, over-forecasting or under-forecasting, and particularly those related to the tails of the predictive distribution [112]. It is negatively oriented and assigns asymmetric weights to negative and positive errors for each quantile.

Definition 3.2.5 (Pinball loss). The pinball loss function for a given quantile q is

$$\rho_q(\hat{x}, x) := \max \left\{ (1 - q)(\hat{x} - x), q(x - \hat{x}) \right\}. \quad (3.11)$$

Definition 3.2.6 (Quantile score). The *quantile score*, over an evaluation set of length N for lead time k , is

$$\text{QS}(k) := \frac{1}{N} \sum_{t=1}^N \frac{1}{Q} \sum_{q=1}^Q \rho_q(\hat{x}_{t+k|t}^{(q)}, x_{t+k}). \quad (3.12)$$

The quantile score, over an evaluation set of length N for a given quantile q over all lead times k , is

$$\text{QS}(q) := \frac{1}{N} \sum_{t=1}^N \frac{1}{T} \sum_{k=k_1}^{k_T} \rho_q(\hat{x}_{t+k|t}^{(q)}, x_{t+k}). \quad (3.13)$$

Interval score

The interval score (IS) [70] assesses the quality of central prediction interval forecasts specifically. It rewards narrow prediction intervals. It penalizes the forecasts where the observation is outside the interval thanks to the penalty term that depends on α

Definition 3.2.7 (IS). The *interval score*, for a central prediction interval with a coverage rate of $(1 - \alpha)$, over an evaluation set of length N and for lead time k is

$$\begin{aligned} \text{IS}(k) = & \frac{1}{N} \sum_{t=1}^N (\hat{x}_{t+k|t}^{(1-\alpha/2)} - \hat{x}_{t+k|t}^{(\alpha/2)}) + \frac{2}{\alpha} (\hat{x}_{t+k|t}^{(\alpha/2)} - x_{t+k}) 1_{\{x_{t+k} \leq \hat{x}_{t+k|t}^{(\alpha/2)}\}} \\ & + \frac{2}{\alpha} (x_{t+k} - \hat{x}_{t+k|t}^{(1-\alpha/2)}) 1_{\{x_{t+k} \geq \hat{x}_{t+k|t}^{(1-\alpha/2)}\}}. \end{aligned} \quad (3.14)$$

Note: the skill scores can be normalized by the total installed capacity for a renewable generation plant.

3.2.3 Multivariate skill scores

Energy score

The energy score (ES) is the most commonly used scoring rule when a finite number of trajectories represents distributions. It is a multivariate generalization of the CRPS and has been formulated and introduced by Gneiting and Raftery [70]. The ES is proper and negatively oriented, *i.e.*, a lower score represents a better forecast. The ES is used as a multivariate scoring rule by Golestaneh et al. [72] to investigate and

analyze the spatio-temporal dependency of PV generations. They emphasize the ES pros and cons. It is capable of evaluating forecasts relying on marginals with correct variances but biased means. Unfortunately, its ability to detect incorrectly specified correlations between the components of the multivariate quantity is somewhat limited.

Definition 3.2.8 (Energy score). Gneiting and Raftery [70] introduced a generalization of the continuous ranked probability score defined in negative orientation as follows

$$\text{ES}(P, \mathbf{x}) = \mathbb{E}_P \|X - \mathbf{x}\| - \frac{1}{2} \mathbb{E}_P \|X - X'\|, \quad (3.15)$$

where X and X' are independent random variables with distribution P and finite first moment, \mathbb{E}_P is the expectation according to the probabilistic distribution P , and $\|\cdot\|$ the Euclidean norm.

The ES is computed following Gneiting et al. [71] over an evaluation set of length N as follows

$$\text{ES} = \frac{1}{N} \sum_{t=1}^M \left[\frac{1}{M} \sum_{i=1}^M \|\hat{\mathbf{x}}_t^i - \mathbf{x}_t\| - \frac{1}{2M^2} \sum_{i,j=1}^M \|\hat{\mathbf{x}}_t^i - \hat{\mathbf{x}}_t^j\| \right]. \quad (3.16)$$

Note: when we consider the marginals of \mathbf{x} , it is easy to recognize that (3.16) is the CRPS.

Variogram score

An alternative class of proper scoring rules based on the geostatistical concept of variograms is proposed by Scheuerer and Hamill [157]. They study the sensitivity of these variogram-based scoring rules to inaccurate predicted means, variances, and correlations. The results indicate that these scores are distinctly more discriminative concerning the correlation structure. Thus, the Variogram score (VS) captures correlations between multivariate components in contrast to the Energy score.

Definition 3.2.9 (Variogram score). For a given day t of the testing set and a T -variate observation $\mathbf{x}_t \in \mathbb{R}^T$, the Variogram score metric of order γ is formally defined as

$$\text{VS}_t = \sum_{k,k'}^T w_{kk'} \left(|x_{t,k} - x_{t,k'}|^\gamma - \mathbb{E}_P |\hat{x}_{t,k} - \hat{x}_{t,k'}|^\gamma \right)^2, \quad (3.17)$$

where $\hat{x}_{t,k}$ and $\hat{x}_{t,k'}$ are the k -th and k' -th components of the random vector $\hat{\mathbf{x}}_t$ distributed according to P for which the γ -th absolute moment exists, and $w_{kk'}$ are

non-negative weights. Given a set of M scenarios $\{\hat{\mathbf{x}}_t^i\}_{i=1}^M$ for this given day t , the forecast variogram $\mathbb{E}_P|\hat{x}_{t,k} - \hat{x}_{t,k'}|^\gamma$ can be approximated $\forall k, k' = 1, \dots, T$ by

$$\mathbb{E}_P|\hat{x}_{t,k} - \hat{x}_{t,k'}|^\gamma \approx \frac{1}{M} \sum_{i=1}^M |\hat{x}_{t,k}^i - \hat{x}_{t,k'}^i|^\gamma. \quad (3.18)$$

Then, the VS is averaged over the testing set of length N

$$\text{VS} = \frac{1}{N} \sum_{t=1}^N \text{VS}_t. \quad (3.19)$$

In this thesis, we evaluate the Variogram score with equal weights across all hours of the day $w_{kk'} = 1$ and using a γ of 0.5, which for most cases provides a good discriminating ability as reported in Scheuerer and Hamill [157].

3.3 Point forecasting evaluation example

This section proposes an example of point forecasting quality evaluation. It is based on an extract of Dumas et al. [54]. In addition, it introduces multi-output weather-based point forecasting by defining the problem formulation and implementation on a real-case study to compute PV and electrical consumption forecasts. They are used by a day-ahead deterministic planner presented in Chapter 9. Note: PV intraday point forecast will also be considered in this thesis using an encoder-decoder architecture developed in Dumas et al. [51]. Chapter 4 details the approach and the results.

3.3.1 Formulation

Point forecasting corresponds to the conditional expectation of the stochastic process for every lead time (2.1). This definition is linked to the so-called *loss function* $L(\hat{x}, x)$. Loss functions assign a penalty to forecast errors as a proxy of the cost of these errors for those making decisions based on such forecasts. There exist various types of loss functions, such as the mean squared error or the pinball loss. A special relevant case of a loss function to compute point forecasts is the mean squared error

$$L_2(\hat{x}, x) = (\hat{x} - x)^2. \quad (3.20)$$

At the time t a point forecast $\hat{x}_{t+k|t}$ for time $t+k$ is the value of the process such that it minimizes the expected loss for the forecast user for all potential realizations of

the process, given our state of knowledge at that time. Therefore, when considering multi-output point forecasts of length T the estimation of the model parameters is performed by solving

$$\theta^* = \arg \min_{\theta} \frac{1}{N} \sum_{t=1}^N \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2, \quad (3.21a)$$

$$\hat{\mathbf{x}}_t \sim g_{\theta}(\mathbf{x}_{<t}, \mathbf{c}_t), \quad (3.21b)$$

given the information set \mathcal{D} of length N , and $\|\cdot\|_2$ the Euclidean norm.

3.3.2 Forecasting models

Two standard forecasting techniques are implemented to forecast the PV production and the consumption, with a dedicated model per variable of interest. The first model uses a Recurrent Neural Network (RNN) of the Keras Python library [33]. The RNN is a Long Short Term Memory (LSTM) with one hidden layer composed of $2 \times n + 1$ neurons with n the number of input features. The second model is a Gradient Boosting Regressor (GBR) of the Scikit-learn Python library [140]. They both use past values of the PV production and consumption, and the weather forecasts provided by the Laboratory of Climatology of the Liège University, based on the MAR regional climate model [59]. It is an atmosphere model designed for meteorological and climatic research, used for a wide range of applications, from km-scale process studies to continental-scale multi-decade simulations. The study Dumas et al. [54] focuses on the real-time control of microgrids based on planning that requires a forecast horizon of a few hours up to a few days. Both models are trained by solving (3.21a) that becomes

$$\hat{\mathbf{x}}_t \sim g_{\theta}(\mathbf{x}_{<t}, \mathbf{c}_t), \quad (3.22a)$$

$$\mathbf{x}_{<t} = [x_t, \dots, x_{t-k_A}], \quad (3.22b)$$

$$\mathbf{c}_t = [\hat{c}_{t+k_1|t}^i, \dots, \hat{c}_{t+k_T|t}^i], \quad (3.22c)$$

with $\hat{\mathbf{x}}_t$ the variable to forecast, *i.e.*, PV, consumption, and \hat{c}_i the forecast of the i^{th} weather variable, *e.g.*, direct solar irradiance, wind speed, air temperature. In the case study considered, the point forecasts are computed each quarter for the subsequent $T = 96$ periods with a resolution $\Delta t = 15$ minutes. The forecasting process is implemented by using a rolling forecast methodology where the learning set is updated every six hours, with a fixed size limited to the week preceding the forecasts, to maintain a reasonable computation time.




Fig. 3.2 PV point forecasts on *June 12, 2019* computed at 06h00 UTC. The PV observations are in black.

3.3.3 Results

Figure 3.2 illustrates the PV point forecasts of the RNN and GBR models compared to the observations (black) on a particular day of the testing set. It is interesting to point out the sudden drop in the PV production around 10 a.m. that the models do not accurately forecast.

The point forecasts are evaluated each quarter over the entire forecasting horizon, the next 96 periods, using the Normalized Mean Absolute Error (NMAE), the Normalized Root Mean Squared Error (NRMSE), and the Normalized Energy Measurement Error (NEME). The normalizing coefficients for both the NMAE and NRMSE are the mean of the absolute values of the PV and consumption over all the simulation data set. The NEME is the NMAE of the energy summed over the entire forecasting horizon. Figures 3.3 and 3.4 provide the scores for both GBR and LSTM models for each quarter, plain lines, and the average over the entire simulation dataset, dashed lines. On average, the LSTM model yields slightly smaller average NMAE, NRMSE, and NEME values. However, the differences are not significant, and the forecast value (see Chapter 9) needs to be evaluated to decide which model is the best.

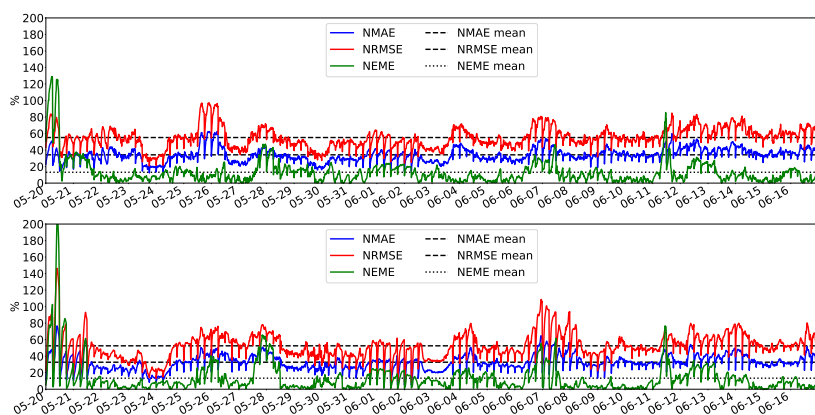


Fig. 3.3 PV forecast scores, GBR (top) and LSTM (bottom).




Fig. 3.4 Consumption forecast scores, GBR (top) and LSTM (bottom).

3.4 Conclusions

This Chapter provides the basics of forecast verification to acquire background knowledge on forecast quality that corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. In contrast, the forecast value, investigated in Part II, relates to the interests of using forecasts in decision-making, such as participation in the electricity market. However, it is intuitively expected that higher-quality forecasts will yield better policies and decisions. Therefore, quality evaluation is complementary to value evaluation and provides an insight into forecasting model skills. The Chapter illustrates the evaluation methodology with a multi-output point forecasting model and an implementation of a real-case study using two standard techniques. The weather-based forecasting models are trained to compute quarterly day-ahead PV and electrical consumption point forecasts. Then, the quality is assessed by computing point forecasts metrics. In the following Chapters, the quality metrics are used to evaluate the various probabilistic forecasts on several case studies.

Chapter 4

Quantile forecasting

Overview

Overall, the Chapter contributions can be summarized as follows.

1. A deep learning-based multi-output quantile architecture computes prediction intervals of PV generation on a day-ahead and intraday basis. The goal is to implement an improved probabilistic intraday forecaster, the encoder-decoder, to benefit from the last PV generation observations. This architecture is compared to a feed-forward neural network.
2. The weather forecasts are used to directly take into account the impact of the weather forecast updates generated every six hours.
3. A proper assessment of the quantile forecasts is conducted by using a k -fold cross-validation methodology and probabilistic metrics. It allows computing average scores over several testing sets and mitigating the dependency of the results to specific days of the dataset.
4. Finally, a comparison of deep learning quantile regression models is conducted with quantiles derived from deep learning generative models.

References: This chapter is an adapted version of the following publications:

Jonathan Dumas, Colin Cointe, Xavier Fettweis, and Bertrand Cornélusse. Deep learning-based multi-output quantile forecasting of pv generation. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494976. URL <https://arxiv.org/abs/2106.01271>.

Jonathan Dumas, Antoine Wehenkel, Damien Lanaspéze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

“Prediction is very difficult, especially if it’s about the future.”

— Niels Bohr

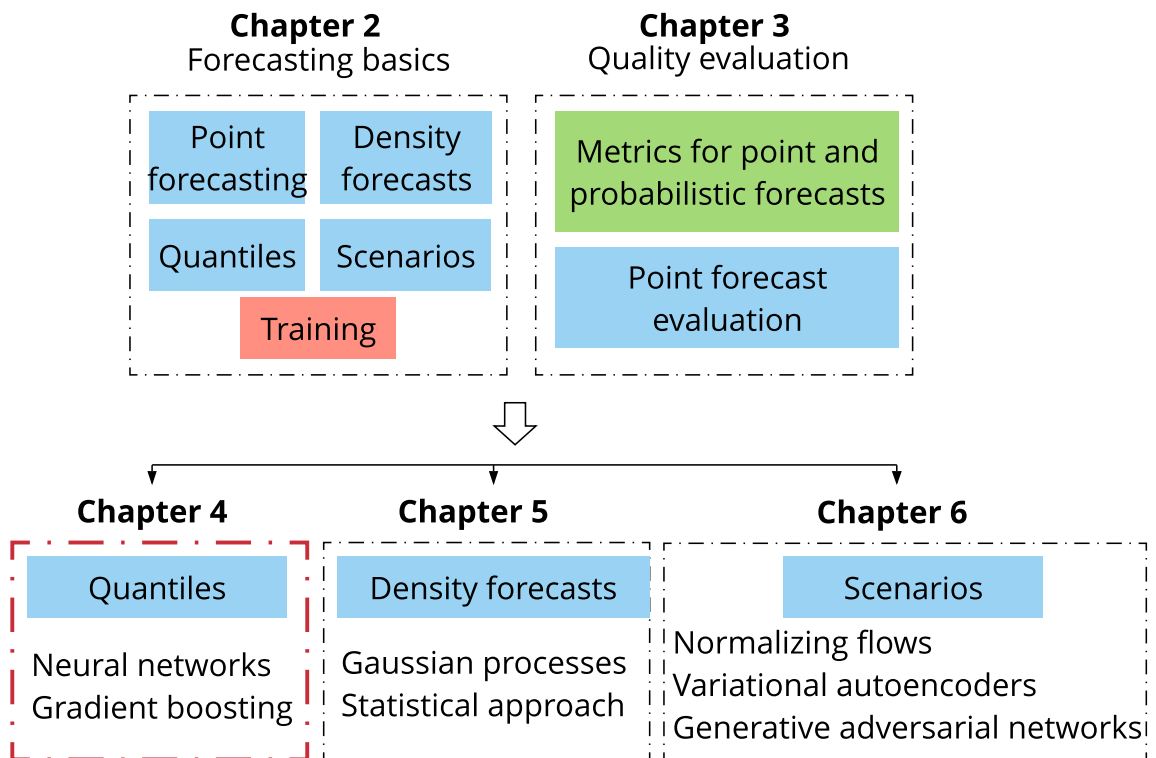


Fig. 4.1 Chapter 4 position in Part I.

This Chapter focuses on *quantile* forecasts that provide probabilistic information about future renewable power generation, in the form of a threshold level associated with a probability. This approach is investigated in Dumas et al. [51] with probabilistic PV forecasters that exploit recent breakthroughs in the field of data science by using advanced deep learning structures, such as the *encoder-decoder* architecture [27]. It is implemented to compute day-ahead and intraday multi-output PV quantiles forecasts, to efficiently capture the correlation between time periods, and to be used as input of a robust optimization model. For instance, to address the energy management system of a grid-connected renewable generation plant coupled with a battery energy storage device detailed in Dumas et al. [52] and Chapter 12. The case study is composed of PV production monitored on-site at the University of Liège (ULiège), Belgium. The weather forecasts from the MAR climate regional model [59] provided by the Laboratory of Climatology are used as inputs of the deep learning models.

The remainder of this Chapter is organized as follows. Section 4.1 presents the non-parametric quantile forecasting framework. Section 4.2 provides the related work. Section 4.3 details the forecasting techniques. Section 4.4 describes the case study and presents the results. Section 4.5 proposes a comparison of the quantile regression models with quantiles derived from deep learning generative models. Finally, Section 4.6 summarizes the main findings and highlights ideas for further work.

4.1 Formulation

Quantile regression [106] is one of the most famous non-parametric approach. It does not assume the shape of the predictive distributions. It can be implemented with various types of forecasting techniques, *e.g.*, neural networks, linear regression, gradient boosting, or any other regression techniques. The estimation of the model parameters θ to compute multi-output quantile forecasts of length T is performed by minimizing the pinball loss (3.11) over all lead times

$$\theta^* = \arg \min_{\theta} \frac{1}{N} \sum_{t=1}^N \frac{1}{Q} \sum_{q=1}^Q \sum_{k=k_1}^{k_T} \rho_q(\hat{x}_{t+k}^{(q)}, x_{t+k}), \quad (4.1)$$

$$\hat{\mathbf{x}}_t^{(q)} \sim g_{\theta}(\mathbf{x}_{<t}, \mathbf{c}_t), \quad (4.2)$$

given the information set \mathcal{D} of length N , and a set of Q quantiles.

4.2 Related work

The literature on quantile forecasting is broad. We present a few papers that have gained our attention in quantile PV probabilistic forecasting. At the Global Energy Forecasting Competition 2014 [85] solar forecasts are expressed in the form of 99 quantiles with various nominal proportions between zero and one. The models are evaluated by using the pinball loss function. This study summarizes the recent research progress on probabilistic energy forecasting at that time, and this competition made it possible to develop innovative techniques. A systematic framework for generating PV probabilistic forecasts is developed by Golestaneh et al. [73]. A non-parametric density forecasting method based on Extreme Learning Machine is adopted to avoid restrictive assumptions on the shape of the forecast densities. A combination of bidirectional Recurrent Neural Networks (RNNs) with Long Short-Term Memory (LSTM), called Bidirectional LSTM (BLSTM), is proposed by Toubeau et al. [166]. It has the benefits of both long-range memory and bidirectional processing. The BLSTM is trained by minimizing the quantile loss to compute quantile forecasts of aggregated load, wind and PV generation, and electricity prices on a day-ahead basis. Finally, an innovative architecture, referred to as encoder-decoder (ED), is developed by Bottieau et al. [27] to generate reliable predictions of the future system imbalance used for robust optimization.

4.3 Forecasting models

This Section presents the forecasting techniques implemented to compute multi-output point and quantile forecasts of PV generation.

4.3.1 Gradient boosting regression (GBR)

Gradient boosting builds an additive model in a forward stage-wise fashion [83]. It allows for the optimization of arbitrary differentiable loss functions. In each stage, a regression tree is fit on the negative gradient of the given loss function. The gradient boosting regressor (GBR) from the Scikit-learn [140] Python library is trained by minimizing the quantile loss. The learning rate is set to 10^{-2} , the max depth to 5, and the number of estimators to 500. There is a GBR model trained per quantile as this library does not handle a single model for several quantiles.

4.3.2 Multi-layer perceptron (MLP)

A description of the most widely used "vanilla" neural network, the *multi-layer perceptron* (MLP), is provided by [83]. A MLP with a single hidden layer is considered for the day-ahead forecasts and as the benchmark for the intraday forecasts. Note, MLPs with two and three hidden layers did not provide any significant improvement on the case study considered. The activation function is the Rectified Linear Unit (ReLU). The number of neurons of the hidden layer is $n_{\text{input}} + (n_{\text{output}} - n_{\text{input}})/2$, with n_{input} and n_{output} the number of neurons of the input and output layers, respectively. The learning rate is set to 10^{-2} and the number of epoch to 500 with a batch size of 8. It is implemented using the PyTorch Python library [139].

4.3.3 Encoder-decoder (ED)

Several technical information about recent advances in neural networks is provided by Toubeau et al. [166], Bottieau et al. [27]. In particular, recurrent neural networks have shown a high potential in processing and predicting complex time series with multi-scale dynamics. However, RNNs are known to struggle in accessing time dependencies more than a few time steps long due to the vanishing gradient problem. Indeed, back-propagated errors during the training stage either fades or blows up over time. Long Short-Term Memory and Gated Recurrent Units networks tackle this problem by using internal memory cells [27]. A neural network composed of a LSTM and feed-forward layers, referred to as LSTM in the rest of the Chapter, is implemented for the day-ahead and intraday forecasts. The number of LSTM units is $n_{\text{input}} + (n_{\text{output}} - n_{\text{input}})/3$, and the number of neurons of the feed-forward layer $n_{\text{input}} + 2 \times (n_{\text{output}} - n_{\text{input}})/3$.

An innovative architecture referred to as encoder-decoder [27], is composed of two different networks and has recently shown promising results for translation tasks and speech recognition applications and imbalance price forecasting. The encoder-decoder, depicted in Figure 4.2, processes features from the past, such as past PV observations, to extract the relevant historical information that is contained into a reduced vector of fixed dimensions, based on the last hidden state. Then, the decoder processes this representation along with the known future information such as weather forecasts. A version of the encoder-decoder architecture (ED-1) is implemented with a LSTM as the encoder and a MLP as the decoder. In a second version (ED-2), the decoder is a LSTM followed by an additional feed-forward layer. Both versions of the encoder-decoder are used as intraday forecasters. In ED-1, the encoder has $2 \times n_{\text{input}}$ units with n_{input} the number of neurons of the encoder input layer, features from the past. Then, the

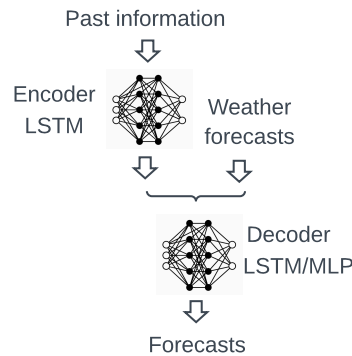


Fig. 4.2 Encoder-decoder architecture.

encoder output is merged with the weather forecasts becoming the decoder input layer that has $n_{\text{output}}/2$ neurons. In ED-2, the decoder has the same number of cells as the encoder, and the feed-forward layer is composed of $n_{\text{output}}/2$ neurons. The LSTM, ED-1, and ED-2 models are implemented using the TensorFlow Python library [1]. The activation functions are the ReLU. The learning rate is set to 10^{-3} , the number of epoch to 500 with a batch size of 64 for the three models.

A sensitivity analysis has been conducted to select the hyperparameters: number of hidden layers, neurons, epochs, and learning rate. Overall, increasing the number of hidden layers and neurons increases the model complexity. It can enhance the accuracy, but only up to a limited number of layers and neurons due to overfitting issues. In addition, the hyperparameter solution is closely related to the size of the historical database [166]. A deep learning model with more significant hidden layers and neurons requires a large amount of data to estimate the parameters accurately. In the case study considered, there are only 157 days of data with a 15 minutes resolution. Thus, we decided to restrict the number of layers and neurons to select a smaller model that performs better with the available information.

4.4 The ULiège case study

4.4.1 Case study description

The ULiège case study is composed of a PV generation plant with an installed capacity of 466.4 kW. The PV generation has been monitored on a minute basis for 157 days. The data is resampled to 15 minutes. The set of quantiles is $\{q = 10\%, \dots, 90\%\}$ for both the day-ahead and intraday forecasts. Numerical experiments are performed on

an Intel Core i7-8700 3.20 GHz based computer with 12 physical CPU cores and 32 GB of RAM running on Ubuntu 18.04 LTS.

4.4.2 Numerical settings

The MAR regional climate model [59] provided by the Laboratory of Climatology of the Liège University is forced by GFS (Global Forecast System) to compute weather forecasts on a six hours basis, four-time gates per day at 00:00, 06:00, 12:00, and 18:00 with a 10-day horizon and a 15 minutes resolution. The solar irradiance and air temperature at 2 meters are normalized by a standard scaler and used as inputs to the forecasting models.

A k -fold cross-validation strategy computes average scores over several testing sets to mitigate the dependency of the results to specific days of the dataset. The dataset is divided into k parts of equal length, and there are k possible testing sets $1 \leq i \leq k$. For a given testing set i , the models are trained over the $k - 1$ parts of the dataset. Eleven pairs of fixed lengths of 142 and 15 days are built. One pair is used to conduct the hyperparameters sensitivity analysis, and the ten others for testing where the scores are averaged. The Mean Absolute Error and Root Mean Squared Error are introduced to evaluate the point forecasts. The MAE, RMSE, CRPS, and IS are normalized by the PV total installed capacity with NMAE and NRMSE the normalized MAE and RMSE.

The day-ahead models, MLP, LSTM, and GBR compute forecasts at noon for the next day. Four intraday time gates are considered at 00:00, 06:00, 12:00, and 18:00. The intraday forecasts of time gate 00:00 are computed by the day-ahead models using only the weather forecasts. Then, the following three intraday forecasts are computed by intraday models where the MLP, ED-1, and ED-2 models use the weather forecasts and the last three hours of PV generation.

The day-ahead and the first intraday predictions are delivered for the 96 quarters of the next day from 00:00 to 23:45 indexed by time steps $0 \leq k \leq 95$. The prediction horizons span from 12 to 36 hours, for the day-ahead gate 12:00, and 0 to 24 hours, for the intraday gate 00:00. The prediction horizon is cropped to $11 \leq k \leq 80$ because the PV generation is always 0 for time steps $0 \leq k \leq 10$ and $81 \leq k \leq 95$ on the ULiège case study. The next three intraday predictions are performed for the 72, 48, and 24 next quarters of the day corresponding to the gates 06:00, 12:00, and 18:00. Therefore, the prediction horizons span from 0 to 18 hours, 0 to 12 hours, and 0 to 6 hours. The intraday forecasting time periods are $24 \leq k \leq 80$, $48 \leq k \leq 80$, and $72 \leq k \leq 80$.

Table 4.1 compares the mean and the standard deviation of the computation times over the ten learning sets to train the point and quantile forecast models¹.

day-ahead	MLP	LSTM	GBR
point	5.3 (0.1)	23.7 (0.3)	3.4 (0.1)
quantile	7.6 (0.2)	69.0 (0.6)	44.6 (0.4)
intraday	MLP	ED-1	ED-2
point	5.0 (0.1)	5.2 (0.1)	17.2 (0.2)
quantile	17.9 (0.2)	6.4 (0.2)	18.0 (0.3)

Table 4.1 Training computation time (s).

4.4.3 Day-ahead results

Figure 4.3a compares the NMAE (plain lines), NRMSE (dashed lines), and Figure 4.3b the CRPS per lead time k of the day-ahead models of gate 12:00. Table 4.2 provides the mean and standard deviation of the NMAE, NRMSE, and CRPS. The LSTM achieved the best results for both point and quantile forecasts. Figures 4.5a, 4.5c, and 4.5e compare the MLP, LSTM, and GBR day-ahead quantile and point forecasts (black line named dad 12) of gate 12:00 on *August 2, 2020* with the observation in red. One can see that the predicted intervals of the LSTM model better encompass the actual realizations of uncertainties than the MLP and GBR.

Score	Gate	MLP	LSTM	GBR
NMAE	12	8.2 (1.2)	7.6 (1.5)	9.2 (0.9)
	24	7.9 (1.2)	7.7 (1.6)	9.0 (0.8)
NRMSE	12	10.2 (1.4)	9.2 (1.6)	11.2 (0.9)
	24	9.7 (1.2)	9.4 (1.8)	10.9 (0.8)
CRPS	12	6.2 (1.1)	4.4 (0.2)	6.4 (0.7)
	24	6.2 (1.0)	4.4 (0.2)	6.3 (0.6)

Table 4.2 The NMAE, NRMSE, CRPS of day-ahead models are averaged over all periods. The best-performing model for each score and gate is written in bold.

¹The day-ahead and intraday LSTM training times are identicals for both point and quantile forecasts as they only take the weather forecasts as inputs.




Fig. 4.3 NMAE, NRMSE, and CRPS of the day-ahead models per lead time k .

4.4.4 Intraday results

Table 4.3 provides the averaged NMAE, NRMSE, and CRPS per gate of intraday models. The LSTM achieved the best NMAE and NRMSE for the 06:00 gate, and the ED-1 achieved the best NMAE and NRMSE for the noon gate and the best CRPS for both gates. Figure 4.4 compares the CRPS per lead time k of the intraday models. The ED-1 benefits from the last PV generation observations. Indeed, some CRPS values for both 06:00 and 12:00 gates are below the ones of 00:00 gate. Table 4.4 provides the interval score of intraday models for 80%, 60%, 40%, and 20% width of central intervals. The ED-1 model achieved the best results except for the prediction interval width of 80%, where it is ED-2 and LSTM for the 06:00 and 12:00 gates, respectively. Overall, The LSTM achieved close results to the ED-1. Figures 4.5b, 4.5d, and 4.5f compare the ED-1, LSTM, and ED-2 intraday quantile and point forecasts (black line named intra 6) of 06:00 gate on *August 2, 2020* with the observation in red. Generally, one can see that the predicted intervals of ED-1 and LSTM models better encompass the actual realizations of uncertainties than ED-2.

Score	Gate	MLP	ED-1	ED-2	LSTM
NMAE	6	8.9 (1.0)	8.5 (1.4)	9.4 (1.0)	7.6 (1.5)
	12	6.7 (1.4)	6.4 (1.3)	7.1 (1.1)	7.2 (1.1)
NRMSE	6	10.9 (0.9)	10.3 (1.3)	11.3 (1.1)	7.7 (1.6)
	12	8.7 (1.3)	7.8 (1.2)	8.5 (1.2)	9.4 (1.8)
CRPS	6	8.1 (0.7)	5.9 (0.9)	6.6 (0.7)	6.2 (0.7)
	12	5.8 (1.2)	4.5 (0.7)	5.6 (1.8)	4.7 (0.5)

Table 4.3 The NMAE, NRMSE, CRPS of intraday models are averaged over all periods. The best-performing model for each score and gate is written in bold.

Width	Gate	MLP	ED-1	ED-2	LSTM
80%	6	24.4 (2.9)	14.9 (4.0)	13.9 (4.9)	19.3 (4.2)
	12	17.4 (3.5)	10.6 (1.8)	11.6 (10.1)	9.6 (2.0)
60%	6	37.6 (3.2)	29.9 (5.0)	32.2 (4.2)	30.7 (4.6)
	12	27.2 (4.3)	22.4 (4.2)	27.5 (10.8)	22.6 (3.1)
40%	6	58.0 (4.5)	50.1 (6.5)	57.2 (6.0)	51.6 (5.8)
	12	42.4 (6.9)	37.7 (5.9)	48.1 (16.8)	39.2 (4.9)
20%	6	111.8 (8.4)	97.1 (11.7)	112.1 (10.3)	99.5 (10.4)
	12	81.5 (13.8)	72.7 (10.0)	94.8 (32.1)	76.5 (8.0)

Table 4.4 The interval score of intraday models is averaged over all periods. The best-performing model for each score and gate is written in bold.




Fig. 4.4 CRPS of intraday models per lead time k .

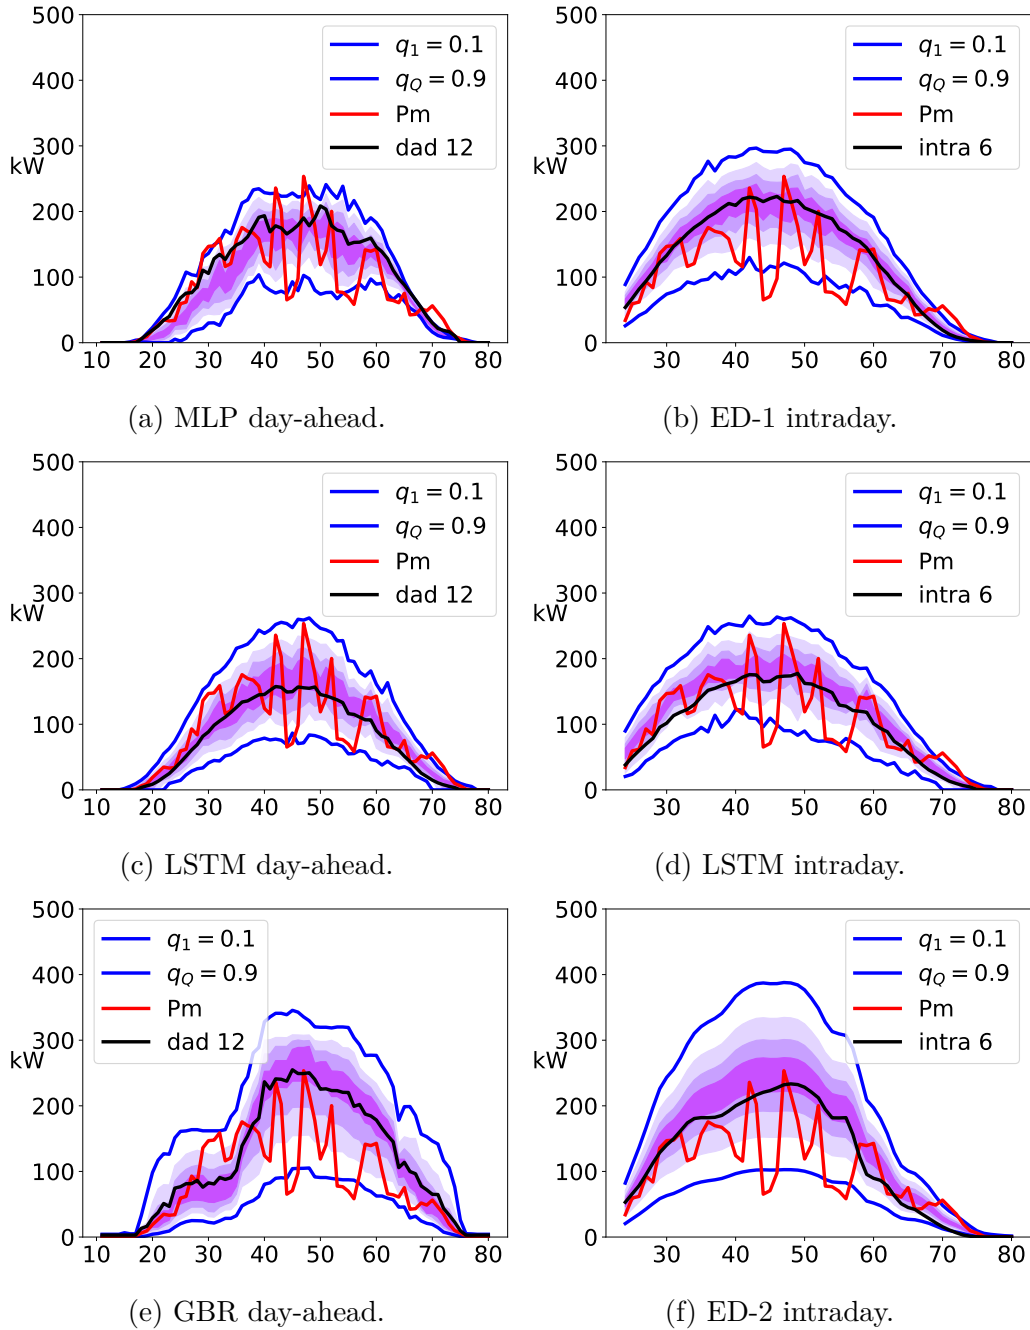


Fig. 4.5 Quantiles *vs.* point forecasts of day-ahead models of gate 12:00 (left), and intraday models of gate 06:00 (right) on *August 2, 2020*, the observations are in red.

4.4.5 Conclusions

An encoder-decoder architecture is implemented on the intraday scale to produce accurate forecasts. It efficiently captures the contextual information composed of past PV observations and future weather forecasts while capturing the temporal dependency between forecasting periods over the entire forecasting horizon. The models are compared by using a k -fold cross-validation methodology and quality metrics on a real case study composed of the PV generation of the parking rooftops of the Liège University. The best day-ahead model for point and quantile forecasts is a neural network composed of a LSTM cell and an additional feed-forward layer. Then, the encoder-architecture composed of a LSTM-MLP yields accurate and calibrated forecast distributions learned from the historical dataset compared to the MLP and LSTM-LSTM models for the intraday point and quantile forecasts. However, the LSTM produced similar results.

Several extensions are under investigation. First, using a larger dataset of at least one full year to consider the entire PV seasonality. Second, developing a PV scenario approach based on the encoder-decoder architecture.

4.5 Comparison with generative models

This Section proposes a quality evaluation of the normalizing flows (NFs) and LSTM PV quantiles, used in Chapter 12 for robust optimization, using the quantile score, the reliability diagram, and the continuous ranked probability score. Indeed, the two-phase engagement control of the capacity firming framework requires day-ahead and intraday *top-quality* forecasts. The more accurate the forecasts, the better the planning and the control. To this end, the Normalizing Flows technique computes quantile day-ahead forecasts compared to a common alternative technique using a Long Short-Term Memory neural network. The controller requires intraday point forecasts computed by an encoder-decoder architecture. NFs are investigated in Section 6.2 of Chapter 6. Both NFs and LSTM models use as input the weather forecasts of the MAR climate regional model provided by the Laboratory of Climatology of the Liège University [59]. The NFs model generates day-ahead scenarios, and the quantiles are derived. The LSTM model computes the quantiles directly and is trained by minimizing the quantile loss. The set of PV quantiles considered for the assessment is $\{q = 10\%, \dots, 90\%\}$.

In this study, the class of Affine Autoregressive flows is implemented². A five-step Affine Autoregressive flow is trained by maximum likelihood estimation with 500 epochs and a learning rate of 10^{-4} . The LSTM learning rate is 10^{-3} , the number of epoch to 500 with a batch size of 64. Figure 4.6 provides the results for these quality metrics

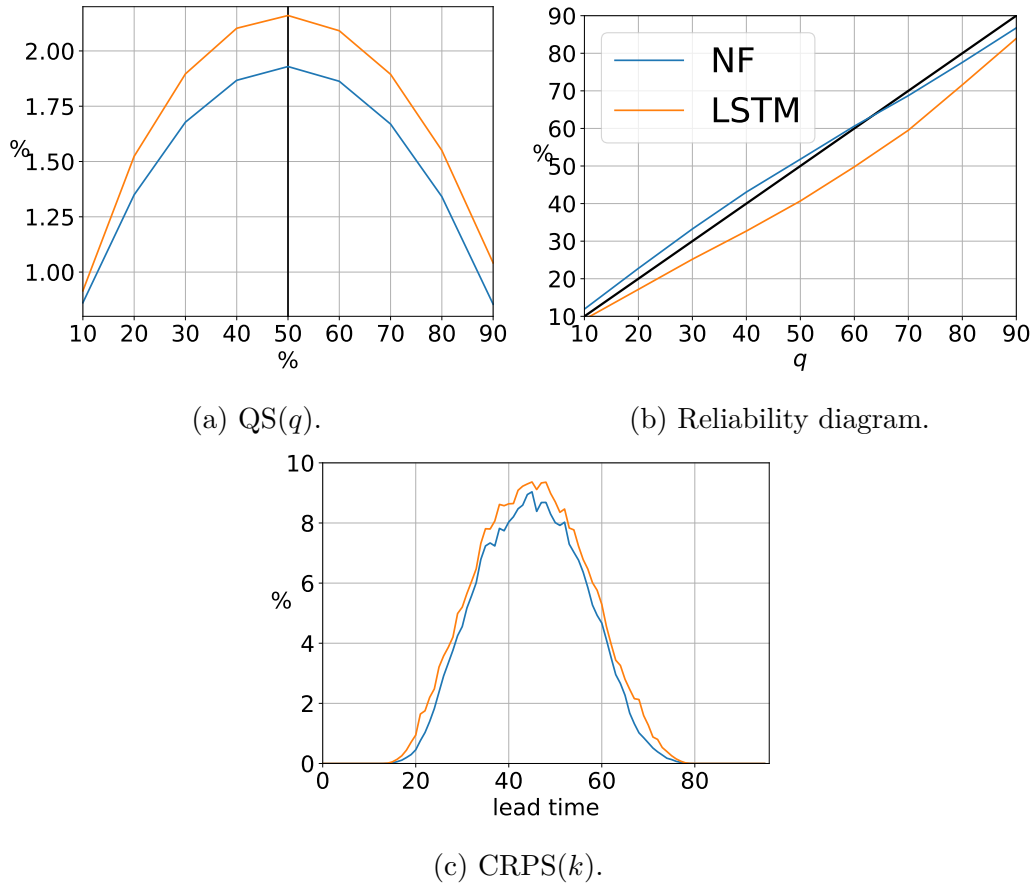


Fig. 4.6 Quantile forecast quality evaluation of LSTM *v.s* NFs models.

computed over the entire dataset normalized by the total installed capacity. The NFs model outperforms the LSTM model with average values of 1.49% and 2.80% *vs.* 1.69% and 3.15% for the QS and CRPS, respectively. The NFs quantiles are also more reliable, as indicated by the reliability diagram. These results motivate the use of the NFs as they outperform common deep learning approaches such as LSTM models.

²<https://github.com/AWehenkel/Normalizing-Flows> [176]

4.6 Conclusions

This Chapter proposes a formulation of the quantile forecasts problem using deep learning models trained by quantile regression to compute multi-output PV quantiles. The forecast quality is evaluated on a real case study composed of the PV generation of the parking rooftops of the Liège University. In addition, these quantile regression models are compared to PV quantiles derived from deep learning generative models, which will be investigated in detail in Chapter 6. In terms of forecast quality, the generative models outperform on this case study the quantile regression models. However, it does not mean they are better in terms of forecast value which will be assessed in Chapter 12 where a robust planner uses the PV quantiles in the form of prediction intervals.

Chapter 5

Density forecasting of imbalance prices

Overview

The contributions of this Chapter are two-fold.

1. A novel two-step probabilistic approach (TSPA) is proposed for forecasting the Belgium imbalance prices. The TSPA uses a direct forecasting strategy [163]. It consists of forecasting an imbalance price for each quarter of the horizon independently from the others, requiring a model per quarter.
2. It sets a reference for other studies as this subject is rarely addressed.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Ioannis Boukas, Miguel Manuel de Villena, Sébastien Mathieu, and Bertrand Cornélusse. Probabilistic forecasting of imbalance prices in the belgian context. In *2019 16th International Conference on the European Energy Market (EEM)*, pages 1–7. IEEE, 2019. URL <https://arxiv.org/abs/2106.07361>.

“*If you don't know where you're going any road will do.*”

— Lewis Carroll

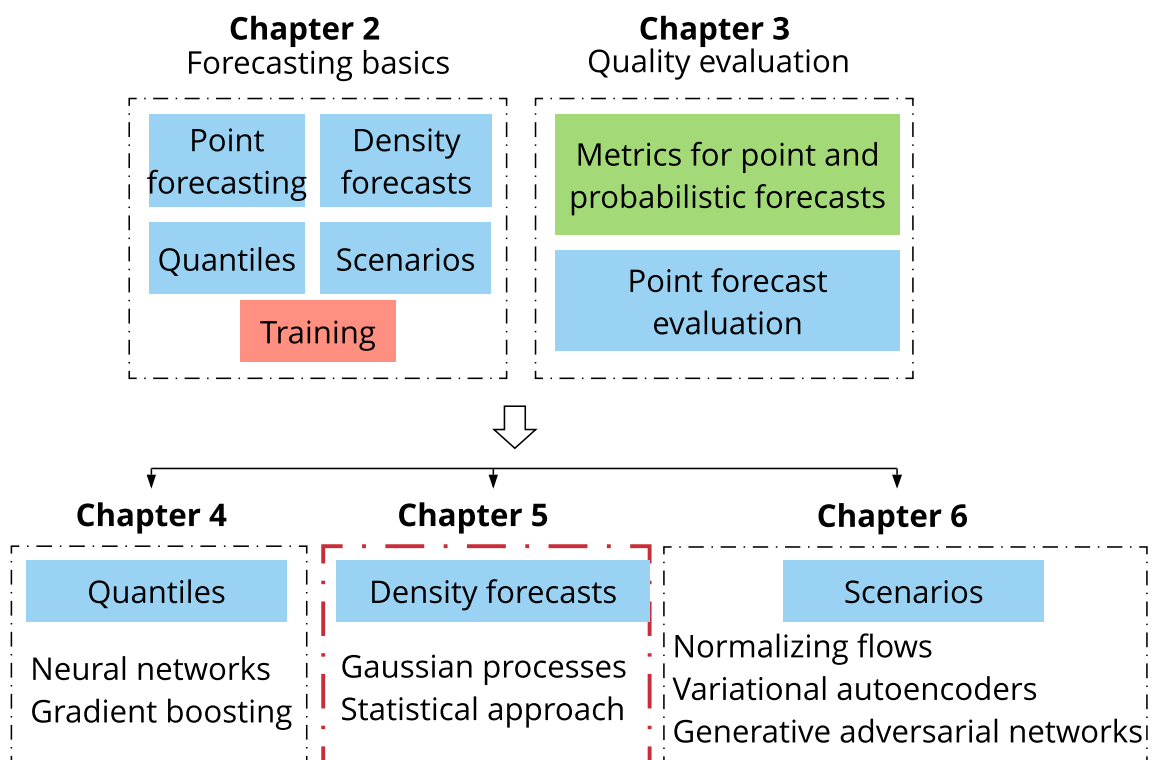


Fig. 5.1 Chapter 5 position in Part I.

This Chapter presents the probabilistic forecasting of imbalance prices methodology developed in Dumas et al. [49]. A novel *two-step probabilistic approach* is proposed, with a particular focus on the Belgian case. The first step consists of computing the net regulation volume (NRV) *state transition probabilities*. It is modeled as a matrix estimated using historical data. This matrix is then used to infer the imbalance prices. Indeed, the NRV can be related to the level of reserves activated, and the corresponding marginal prices for each activation level are published by the Belgian Transmission System Operator (TSO) one day before electricity delivery. The model is compared to: (1) a multi-layer perceptron, implemented in a deterministic setting; (2) the widely used probabilistic technique, the Gaussian Processes (GP).

This Chapter is organized as follows. Section 5.1 details the related work. Section 5.2 introduces the novel two-step probabilistic approach formulation and the assumptions made. Section 5.3 describes the numerical tests on the Belgian case. Section 5.4 reports the results. Conclusions are drawn in Section 5.5. Appendix 5.5 lists the acronyms, parameters, and forecasted or computed variables. Finally, Annex 5.6 provides a short reminder of the imbalance market and the Belgian balancing mechanisms.

5.1 Related work

The progressive, large-scale integration of renewable energy sources has altered electricity market behavior and increased the electricity price volatility over the last few years [41, 77, 99]. In this context, imbalance price forecasting is an essential tool for strategic participation in short-term energy markets. Several studies take into account the imbalance prices as penalties for deviation from the bids to compute the optimal bidding strategy [67, 143, 23, 25]. However, these penalties are known only *a posteriori*. A forecast indicating the imbalance prices and the system position, short or long, with an interval to inform about the range of potential outcomes is a powerful tool for decision making. Probabilistic forecasting usually outperforms deterministic models when used with the appropriate bidding strategies [143]. Whereas the literature on day-ahead electricity forecast models is extensive, studies about balancing market prices forecast have received less attention. We recommend three papers related to this topic. First, a statistical description of imbalance prices for shortage and surplus is presented by Saint-Drenan [154]. Second, a combination of classical and data mining techniques to forecast the system imbalance volume is addressed by Garcia and Kirschen [62]. Finally, a review and benchmark of time series-based methods for balancing market price forecasting are proposed by Klæboe et al. [104]. One-hour and

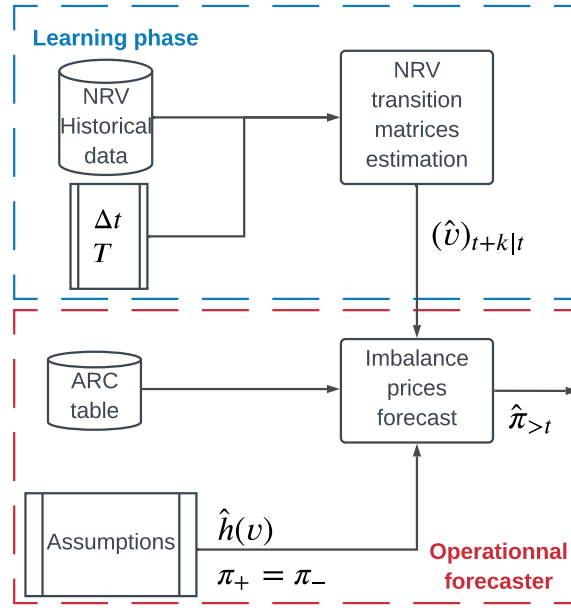


Fig. 5.2 Two-step probabilistic approach (TSPA) imbalance price forecasting process.

one-day-ahead forecasts are considered for state determination, balancing volume, and price forecasting on the Nord Pool price zone NO2 in Norway.

5.2 Formulation

This study focuses on the intraday market time scale that requires a forecasting horizon from a few minutes to a few hours with a resolution Δt . The day-ahead time scale requires forecasts of the imbalance prices from 12 to 36 hours, which is not realistic at this stage. The input data are the imbalance price history, the NRV, and the marginal prices for activation published by the TSO. The probabilistic approach consists of forecasting the imbalance prices in two steps: computing the NRV state transition probabilities, then forecasting the imbalance prices, as depicted in Figure 5.2. It is motivated by the ELIA imbalance price mechanisms detailed in Appendix 5.6.

5.2.1 Net regulation volume forecasting

Let consider the T forecasting horizons $k_1 = \Delta t, \dots, k_T = T\Delta t$ with Δt the market period, 15 minutes for Belgium. The NRV historical data is discretized into N bins, v_i , centered around $v_{i,1/2}$. Note: this discretization has been determined after a statistical study of the NRV distribution. The T NRV transition matrices $(v)_{t+k|t}$, of dimensions $N \times N$, from a known state at time t to a future state at time $t + k$ are estimated by




Fig. 5.3 NRV transition matrix from t to $t + 15$ min (left) and $t + 60$ min (right).

using the NRV historical data, and referred to as $(\hat{v})_{t+k|t}$. They are composed of the following conditional probabilities $\forall k = k_1, \dots, k_T$

$$p_{t+k|t}^{ij} = \Pr[v(t+k) \in v_j \mid v(t) \in v_i], \quad i, j \in \llbracket 1; N \rrbracket^2, \quad (5.1)$$

with $v(t)$ the measured NRV at time t , and $\sum_{j=1}^N p_{t+k|t}^{ij} = 1 \forall i \in \llbracket 1; N \rrbracket$. The conditional probabilities (5.1) are estimated statistically over the learning set (LS) $\forall k = k_1, \dots, k_T$

$$\hat{p}_{t+k|t}^{ij} = \frac{\sum_{t \in LS} 1_{\{v(t+k) \in v_j \mid v(t) \in v_i\}}}{\sum_{t \in LS} 1_{\{v(t) \in v_i\}}}, \quad i, j \in \llbracket 1; N \rrbracket^2. \quad (5.2)$$

Figure 5.3 illustrates the matrices $(\hat{v})_{t+k_1|t}$ and $(\hat{v})_{t+k_4|t}$ with 2017 as learning set. The estimated mean $\hat{v}_{t+k|t}^m$ and standard deviation $\hat{v}_{t+k|t}^{std}$ of the NRV at time t for $t+k$ are calculated as follows

$$\begin{aligned} \hat{v}_{t+k|t}^m &= \sum_{j=1}^N \hat{p}_{t+k|t}^{ij} v_{j,1/2} \\ \hat{v}_{t+k|t}^{std} &= \sqrt{\sum_{j=1}^N \hat{p}_{t+k|t}^{ij} (v_{j,1/2} - \hat{v}_{t+k|t}^m)^2}, \end{aligned} \quad (5.3)$$

with i such as $v(t) \in v_i$.

5.2.2 Imbalance price forecasting

The NRV can be related to the level of reserves activated and the corresponding marginal prices for each activation level, published by the TSO one day before electricity delivery. We thus first forecast the NRV and its spread among the Gross Upward regulation Volume (GUV) and Gross Downward regulation Volume (GDV). Then, we forecast

the reserve products activated (contracted or not) to select the most probable MIP and MDP into the ARC table. Finally, the mean and the standard deviation of the imbalance price forecast are derived.

However, the ARC table contains only the contracted reserve products. Most of the time, the first activated reserve products come from the non contracted International Grid Control Cooperation platform (IGCC-/+), the contracted secondary reserve (R2-/+) and the non contracted regulation reserves (Bids-/+)¹. For instance, consider a quarter of an hour with an NRV of 150 MW, spread into 170 MW of GUV and 20 MW of GDV. Suppose ELIA activated 80 MW of IGCC+ and 90 MW of R2+. Then, the MIP is given in the marginal activation price of R2+ in the ARC table at the range $[0, 100]$ MW. Suppose that ELIA has activated 20 MW of IGCC+, 20 MW of R2 +, and 130 MW of Bids+. Then, the MIP is given in the marginal activation price of Bids+. However, this is not a contracted reserve, and its price is not in the ARC table. Then, it is more complicated to predict the MIP and consequently the imbalance prices. Therefore, we introduce several simplifying assumptions justified by a statistical study on the 2017 ELIA imbalance data.

Assumption 1. *The NRV is entirely spread into the GUV if the NRV is positive or GDV if the NRV is negative.*

The mean and standard deviation of the GUV and GDV are 109 ± 82 MW *vs.* 17 ± 27 MW when the NRV is positive, while it is 13 ± 20 MW *vs.* 110 ± 73 MW when the NRV is negative. This assumption enables to select directly in the ARC table the marginal price for activation corresponding to the range of activation equal to the NRV, minus IGCC.

Assumption 2. *We do not consider the Bids reserve product. Thus, we assume the NRV spreads over the IGCC and reserve products of the ARC table.*

The percentage of Bids reserve product, positive or negative, activated over each quarter of 2017 is 11.5%.

Assumption 3. *The level of activated IGCC reserve product is modeled by a function \hat{h} of the NRV.*

\hat{h} assigns for a given value of NRV a range of activation p into the ARC table. c_t^p is the ARC marginal price at t and for the activation range p , with $p \in \llbracket 1; P \rrbracket$. If $\hat{h}(v)$ falls

¹More information about the reserve products is available at <http://www.elia.be>.

into the activation range p , then $c_t^p(\hat{h}(v))$ is equal to c_t^p . Due to the 2017 statistical distribution of the IGCC versus the NRV, \hat{h} is defined as follows

$$\hat{h}(x) = \begin{cases} x & \text{if } |x| \leq 100, \\ x - 100 & \text{if } x > 100, \\ x + 100 & \text{if } x < -100. \end{cases} \quad (5.4)$$

The mean and standard deviation (MW) of the IGCC+ and IGCC- are

$$\begin{cases} 17 \pm 25 \ \& \ 23 \pm 24 & \text{if } |NRV| \leq 100, \\ 50 \pm 48 \ \& \ 5 \pm 15 & \text{if } NRV > 100, \\ 2 \pm 10 \ \& \ 67 \pm 47 & \text{if } NRV < -100. \end{cases}$$

Generally, ELIA first tries to activate the IGCC product to balance the system. However, when the system imbalance is too high other reserve products are required.

Assumption 4. *The positive imbalance price is equal to the negative one.*

The positive and negative imbalance means prices are 42.23 and 43.04 €/MWh. They are different 30.38% of the time, but the NMAE and NRMSE are 0.02 and 0.06%. Indeed, the positive and negative prices differ only by a small correction parameter if the system imbalance is greater than 140 MW, cf. Appendix 5.6.2.

Under these assumptions, the estimated mean $\hat{\pi}_{t+k|t}^m$ and standard deviation $\hat{\pi}_{t+k|t}^{std}$ of the imbalance prices at time t for $t+k$ are calculated as follows

$$\begin{aligned} \hat{\pi}_{t+k|t}^m &= \sum_{j=1}^N \hat{p}_{t+k|t}^{ij} c_j^{t+k} (\hat{h}(v_{j,1/2})) \\ \hat{\pi}_{t+k|t}^{std} &= \sqrt{\sum_{j=1}^N \hat{p}_{t+k|t}^{ij} (c_j^{t+k} (\hat{h}(v_{j,1/2})) - \hat{\pi}_{t+k|t}^m)^2}, \end{aligned} \quad (5.5)$$

with i such as $v(t) \in v_i$. Finally, on a quarterly basis a forecast is issued at time t and composed of a set of T couples $\hat{\pi}_{>t} := \left\{ (\hat{\pi}_{t+k|t}^m, \hat{\pi}_{t+k|t}^{std}) \right\}_{k=k_1}^{k_T}$.

5.3 Case study

This approach is compared to a widely used probabilistic technique, the Gaussian Processes, and a "classic" deterministic technique, a Multi-Layer Perceptron (MLP). Both techniques are implemented using the Scikit-learn Python library [140]. The




Fig. 5.4 Rolling forecast strategy.

GP uses Matérn, constant and white noise kernels. The MLP has one hidden layer composed of $2 \times n + 1$ neurons with n input features. The dataset comprises the 2017 and 2018 historical Belgium imbalance price and NRV, available on Elia's website. Both the MLP and GP models forecast the imbalance prices based on the previous twenty-four hours of NRV and imbalance prices, representing a total of 2×96 input features. The MLP uses a Multi-Input Multi-Output (MIMO) strategy, and the GP a Direct strategy [163]². The Direct strategy consists of training a model \hat{g}_k per market period

$$\hat{\pi}_{t+k|t} = \hat{g}_k(\pi_t, \dots, \pi_{t-k_T}, v_t, \dots, v_{t-k_T}), \quad \forall k = k_1, \dots, k_T, \quad (5.6)$$

and the forecast is composed of the T predicted values computed by the T models \hat{g}_k . In contrast, the MIMO strategy consists of training only one model \hat{g} to directly compute the T values of the variable of interest

$$[\hat{\pi}_{t+k_1|t}, \dots, \hat{\pi}_{t+k_T|t}]^T = \hat{g}(\pi_t, \dots, \pi_{t-k_T}, v_t, \dots, v_{t-k_T}). \quad (5.7)$$

For both MIMO and Direct strategies, the forecast is computed quarterly and composed of T values. The forecasting process uses a *rolling forecast strategy* where the training set is updated every month, depicted in Figure 5.4. Its size increases by one month for both the MLP and TSPA techniques, with the initial training set being 2017. However, it is limited to the month preceding the forecast for the GP technique to maintain a reasonable computation time. The validation set is 2018, where each month is forecasted by a model trained on a different learning set.

5.4 Results

The probabilistic forecasts are evaluated using the Pinball Loss Function (PLF) and the Continuous Ranked Probability Score. They are compared to the deterministic

²GP regression with multiple outputs is non-trivial and still a field of active research [172, 117].

k	Technique	NMAE	NRMSE	PLF	CRPS
15 min	MLP	52.74	84.37	-	-
	GP	61.33	98.59	16.48	32.64
	TSPA	61.91	101.24	16.07	31.84
60 min	MLP	61.85	97.26	-	-
	GP	62.13	101.14	16.09	31.87
	TSPA	66.47	105.43	15.22	30.15
360 min	MLP	72.64	112.90	-	-
	GP	72.61	114.56	14.79	29.29
	TSPA	73.35	114.2	14.2	28.12

Table 5.1 Average scores over all lead times k . The best-performing model for each lead time is written in bold.

metrics with the Normalized Mean Absolute Error (NMAE), and the Normalized Root Mean Squared Error (NRMSE) of the mean predicted imbalance prices. The scores $NMAE(k)$, $NRMSE(k)$, $PLF(k)$, and $CRPS(k)$ for a lead time k are computed over the entire validation set. The normalizing coefficient for both the NMAE and NRMSE is 55.02 €/MWh, the mean absolute value of the imbalance prices over 2018.

The forecaster computes the mean and standard deviation of the imbalance prices per market period. Then, a Gaussian distribution generates samples, and the percentiles 1, . . . , 99 are derived. They are used to compute the PLF. The CRPS is computed using the `crps_gaussian` function of the Python `proprscoring`³ library. Table 5.1 presents the average scores over all lead times k for the horizons of 15, 60 and 360 minutes. Figure 5.5a provides the average scores over all lead times k for each forecasting horizon, and Figure 5.5b depicts the score per lead time k for the forecasting horizon of 360 minutes.

Two days, depicted in Figure 5.6, from the validation set are selected to illustrate the results. On *January 8*, 2019, the ELIA system was short on average, leading to a high NRV and imbalance prices. On *January 10*, 2019, the ELIA system was alternatively short and long leading to fluctuating NRV and imbalance prices. The 15 minutes horizon forecasts are depicted in Figure 5.7, where only the last forecasted value for each quarter is shown. The 60 and 360 minutes horizon forecasts are depicted in Figure 5.8 in Appendix 5.6.3. On *January 8*, 2019 the GP provides better results on

³<https://pypi.org/project/proprscoring/>

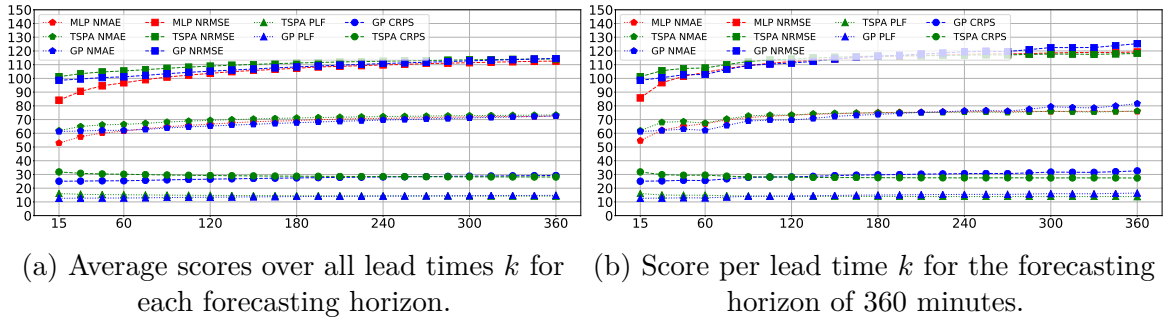


Fig. 5.5 Scores.




Fig. 5.6 ELIA NRV on *January 8, 2019* (blue) and *January 10, 2019* (orange).

average as it accurately follows the actual prices. On *January 10, 2019*, there is no clear winner. Other figures are reported in Appendix 5.6.3 for other forecasting horizons. The MLP provides the best NMAE and NRMSE, except for the horizon of 360 minutes, and the TSPA the best CRPS and PLF scores for the three forecasting horizons considered. However, to select the best forecasting model, it would be necessary to measure the accuracy of the global bidding chain composed of the forecasting and decision-making modules.

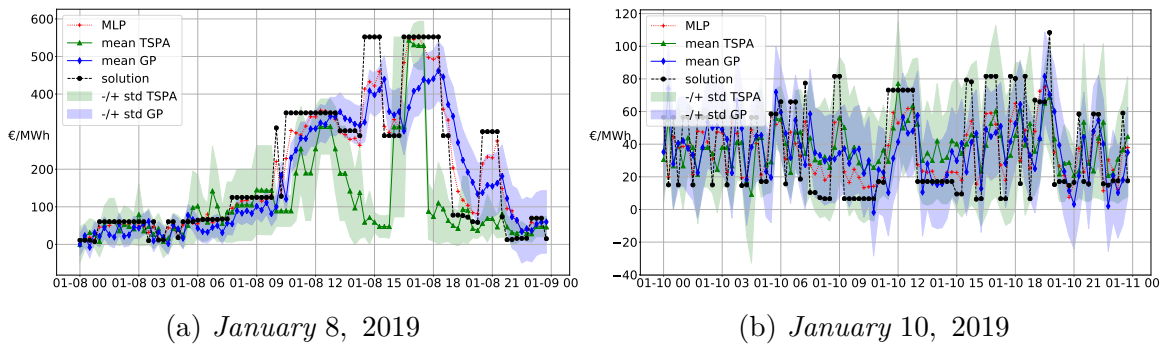


Fig. 5.7 MLP, GP and TSPA 15 minutes horizon forecasts.

5.5 Conclusions

This study addressed the problem of forecasting the imbalance prices in a probabilistic framework. The novel two-step probabilistic approach consists of computing the net regulation volume state transition probabilities. Then, it infers the imbalance price from the ELIA ARC table and computes a probabilistic forecast. It is compared to the MLP and GP techniques in the Belgium case. The results indicate it outperforms them using probabilistic metrics, but it is less accurate at predicting the precise imbalance prices.

Learning models could improve this novel probabilistic approach in three directions: (1) To avoid making our simplifying assumptions. (2) By adding input features to describe the market situation better. (3) Extending the approach to implement the whole bidding strategy chain would allow determining which approach is the best.

Appendix: notation

Acronyms

ARC	Available Regulation Capacity
BRP	Balancing Responsible Party
GP	Gaussian Processes
GDV	Gross Downward regulation Volume
GUV	Gross Upward regulation Volume
IGCC	International Grid Control Cooperation
MDP	Marginal price for Downward Regulation
metric	NMAE, NRMSE, PLF, CRPS
MIMO	Multi-Input Multi-Output
MIP	Marginal price for Upward Regulation
NRV	Net Regulation Volume
R2	Secondary reserve, upwards or downwards
TSO	Transmission System Operator
TSPA	Two-Step Probabilistic Approach

Parameters

Symbol	Description	Unit
π_+, π_-	Positive/Negative imbalance price	€/MWh
α_1, α_2	ELIA parameters for π_+ and π_-	€/MWh
c_t^p	ARC marginal price at t and for activation range p	€/MWh
$v(t)$	NRV measured at time t	MW
v_i	NRV bin i	MW
$v_{i,1/2}$	Center of NRV bin i	MW
$(v)_{t+k t}$	NRV transition matrix from t to $t+k$	-
$p_{t+k t}^{ij}$	NRV conditional probabilities at t for $t+k$	-

Forecasted or computed variables

Symbol	Description	Unit
$\hat{\pi}_{t+k t}^m$	Predicted mean imbalance price at t for $t+k$	€/MWh
$\hat{\pi}_{t+k t}^{std}$	Standard deviation of $\hat{\pi}_{t+k t}^m$ at t for $t+k$	€/MWh
$\hat{\pi}_{>t}$	Set $\left\{ (\hat{\pi}_{t+k t}^m, \hat{\pi}_{t+k t}^{std}) \right\}_{k=k_1}^{k_T}$	€/MWh

$\hat{v}_{t+k t}^m$	Predicted mean NRV at t for $t+k$	MW
$\hat{v}_{t+k t}^{std}$	Standard deviation of $\hat{v}_{t+k t}^m$ at t for $t+k$	MW
$(\hat{v})_{t+k t}$	Estimated NRV transition matrix from t to $t+k$	-
$(\hat{v})_{>t}$	Set $\left\{ (\hat{v})_{t+k t} \right\}_{k=k_1}^{k_T}$	-
$\hat{p}_{t+k t}^{ij}$	Estimated NRV conditional probability at t for $t+k$	-

5.6 Appendix: balancing mechanisms

5.6.1 Balancing mechanisms

A balancing mechanism aims to balance a given geographical area and control sudden imbalances between injection and off-take. Generally, this mechanism relies on exchanges with neighboring TSOs, the responsible balance parties, and the usage of reserve capacities. Each party that desires to inject or off-take to the grid must be managed by a Balancing Responsible Party (BRP). The BRP is responsible for balancing all off-takes and injections within its customer's portfolio. The TSO applies an imbalance tariff when it identifies an imbalance between total physical injections, imports, and purchases on the one hand and total off-takes, exports, and sales on the other. When the BRPs cannot balance their customer's portfolios, the TSO activates reserves to balance the control area. These reserves are mainly from conventional power plants, which can be quickly activated upward or downward to cover real-time system imbalances. The main types of reserve are the Frequency Containment Reserve (FCR), the Automatic Frequency Restoration Reserve (aFRR), the Manual Frequency Restoration Reserve (mFRR), and the Replacement Reserve (RR)⁴. The activation of these reserves results from a merit order representing the activation cost of reserve capacity. If the system faces a power shortage, the TSO activates upward reserves that result in a positive marginal price on the reserve market. Then, the TSO pays the Balancing Service Provider. The cost of this activation is transferred to the BRPs. BRPs facing short positions are reinforcing the system imbalance. They must pay the marginal price to the TSO. BRPs facing long positions are restoring the system imbalance. They receive the marginal price from the TSO. This mechanism incentivizes market players to maintain their portfolios in balance, as well as to reduce the net system imbalance.

⁴<https://www.entsoe.eu/>

5.6.2 Belgium balancing mechanisms

This section describes the ELIA imbalance price mechanisms and the data publication part of the TSPA inputs. On a 15 minutes basis, the NRV is defined as the sum of the GUV and GDV. The Gross Upward Volume (GUV) is the sum of the volumes of all upward regulations. The Gross Downward Volume (GDV) is the sum of the volumes of all downward regulations. If the NRV is positive, the highest price of all upward activated products, the Marginal price for Upward Regulation (MIP), is applied for the imbalance price calculation. If the NRV is negative, the lowest price of all downward activated products, the Marginal price for Downward Regulation (MDP), is applied. The definitions of the positive π_+ and negative π_- imbalance prices are provided in Table 5.2. The correction parameters α_1 and α_2 are zero when the system imbalance is lower than 140 MW and proportional to it when greater than 140 MW.

The MIP and MDP prices are in the third Available Regulation Capacity (ARC) table most of the time. The ARC publication considers the applicable merit order, *i.e.*, the order in which Elia must activate the reserve products. Then, the volumes are ranked by activation price (cheapest first). The marginal price is the highest price for every extra MW upward volume and the lowest for every extra MW downward volume. The ARC table, showing the activation price of the contracted reserves per activation range of 100 MW, displays the estimated activation price considering a certain NRV. For a given quarter-hour t there are P marginal prices for activation c_t^p , $p \in \llbracket 1; P \rrbracket$, each one of them related to the activation range p . P is equal to 22 with 11 negatives ranges and 11 positives ranges. The first activation range, $p = 1$, corresponds to the interval $[-\infty, -1000]$ MW, the second one to $[-1000, -900]$, ..., $[-100, 0]$, $[0, 100]$ up to $[1000, +\infty]$. The data of day D are published on $D - 1$ at 6 pm based on the nomination of day-ahead and intraday programs and bids submitted by the concerned parties. The values of each quarter-hour of the day are refreshed quarterly. Therefore, the published values are an estimation. However, they are likely to include the MIP and MDP prices. At the condition to determine the NRV and its spread between the GUV and GDV. The TSPA takes as input the third ARC table to determine the most probable MIP and MDP prices.

5.6.3 Additional results

BRP perimeter	$NRV < 0$	$NRV > 0$
> 0	$\pi_+ = MDP - \alpha_1$	$\pi_+ = MIP$
< 0	$\pi_- = MDP$	$\pi_- = MIP + \alpha_2$

Table 5.2 Elia imbalance prices.

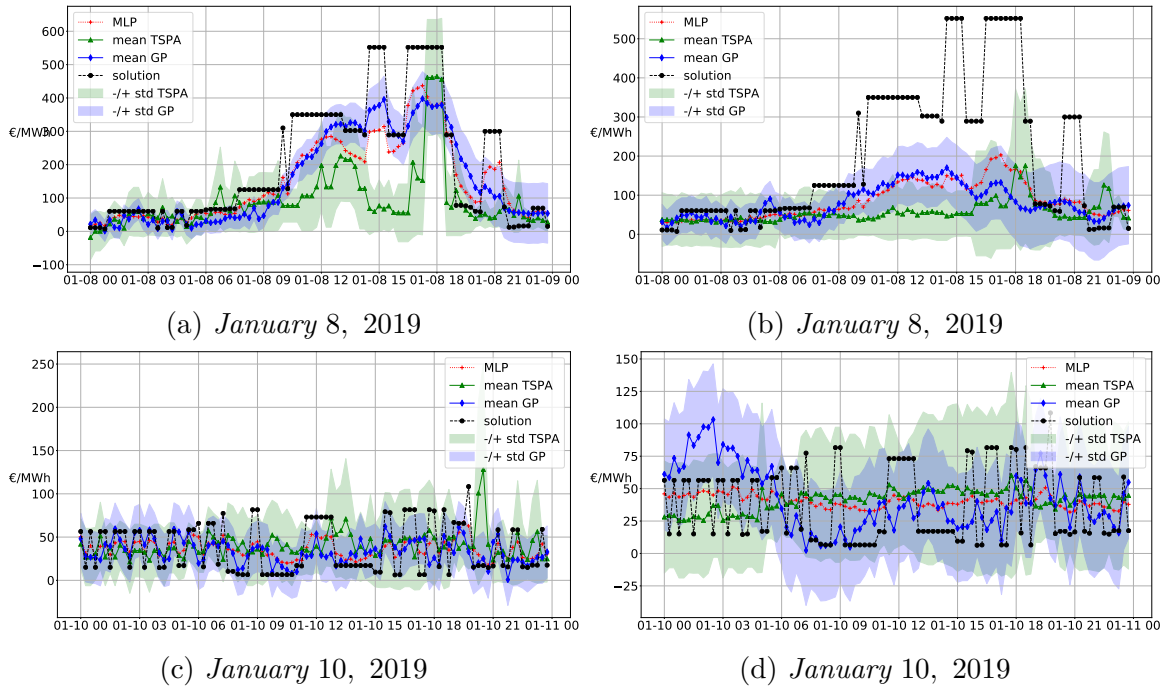


Fig. 5.8 MLP, GP and TSPA 60 (left) and 360 (right) minutes horizon forecasts.

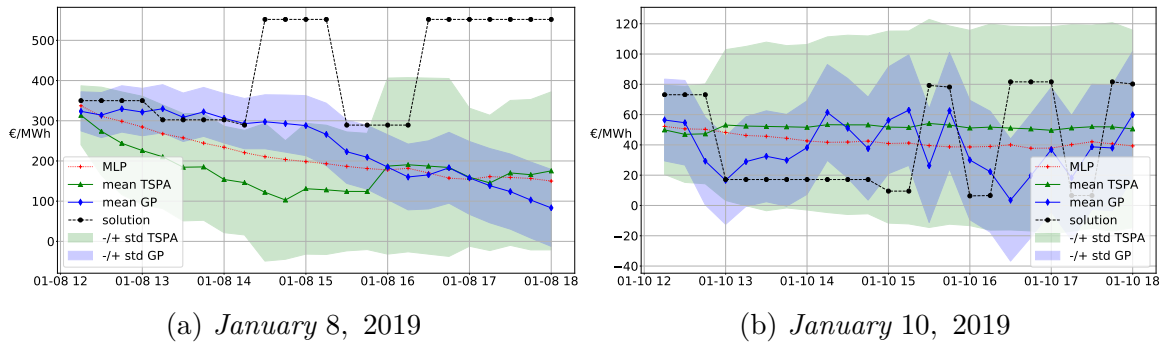


Fig. 5.9 MLP, GP and TSPA forecasts, 12h00 UTC, with an horizon of 360 minutes.

Chapter 6

Scenarios

Overview

The main contributions of this Chapter are three-fold:

1. We present to the power systems community a recent class of deep learning generative models: the Normalizing Flows (NFs). Then, we provide a fair comparison of this technique with the state-of-the-art deep learning generative models, Generative Adversarial Networks (GANs) and Variational AutoEncoders (VAEs). It uses the open data of the Global Energy Forecasting Competition 2014 (GEFcom 2014) [86]. To the best of our knowledge, it is the first study that extensively compares NFs, GANs, and VAEs on several datasets, PV generation, wind generation, and load with a proper assessment of the quality and value based on complementary metrics, and an easily reproducible case study;
2. We implement conditional generative models to compute improved weather-based PV, wind power, and load scenarios. In contrast to most of the previous studies that focused mainly on past observations;
3. Overall, we demonstrate that NFs are more accurate in quality and value, providing further evidence for deep learning practitioners to implement this approach in more advanced power system applications.

This study provides open-access to the **Python code**: <https://github.com/jonathandumas/generative-models>.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Antoine Wehenkel, Damien Lanaspeze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

“ I won't say 'See you tomorrow' because that would be like predicting the future, and I'm pretty sure I can't do that.”

— Ludwig Wittgenstein

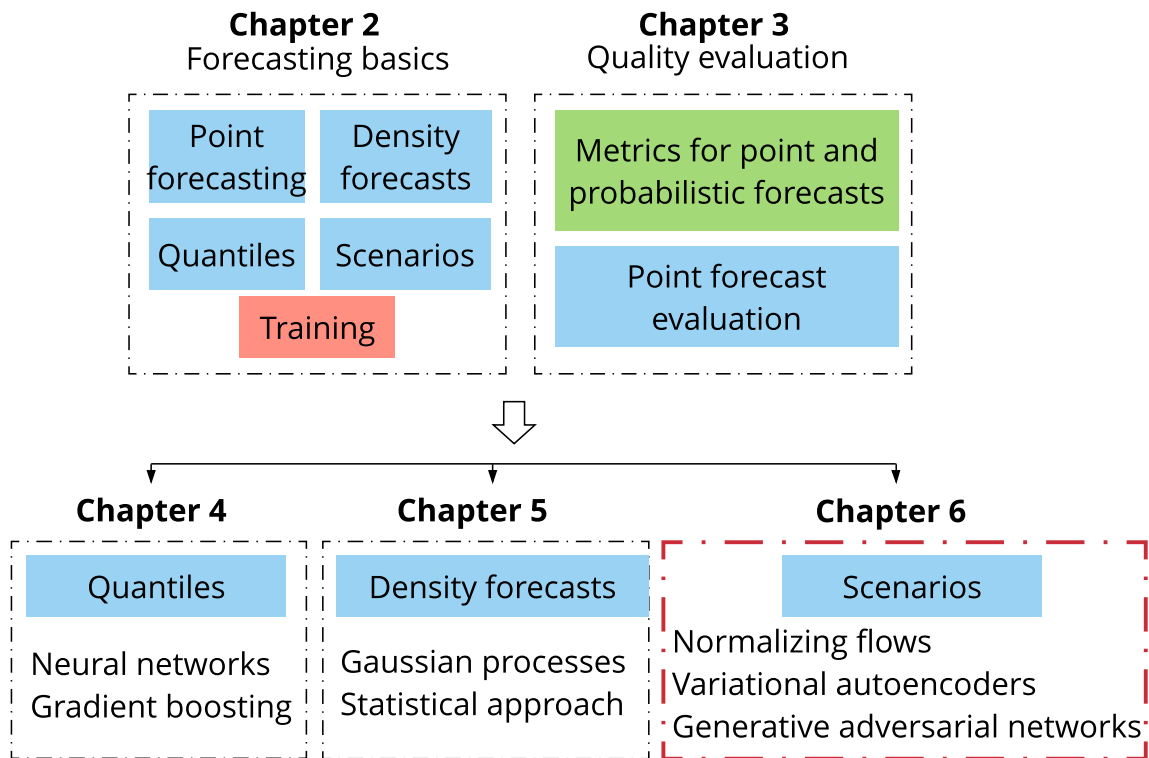


Fig. 6.1 Chapter 6 position in Part I.

6.1 Introduction

This Chapter focuses on *scenario generation*, a popular probabilistic forecasting method to capture the uncertainty of load, photovoltaic (PV) generation, and wind generation. It consists of producing sequences of possible load or power generation realizations for one or more locations.

Forecasting methodologies can typically be classified into two groups: *statistical* and *machine learning* models. On the one hand, statistical approaches are more interpretable than machine learning techniques, sometimes referred to as black-box models. On the other hand, they are generally more robust, user-friendly, and successful in addressing the non-linearity in the data than statistical techniques. We provide in the following a few examples of statistical approaches. More references can be found in Khoshrou and Pauwels [100] and Mashlakov et al. [121].

Multiple linear regression models [173] and autoregressive integrated moving average [40] are among the most fundamental and widely-used models. The latter generates spatiotemporal scenarios with given power generation profiles at each renewables generation site [123]. These models mostly learn a relationship between several explanatory variables and a dependent target variable. However, they require some expert knowledge to formulate the relevant interaction between different variables. Therefore, the performance of such models is only satisfactory if the dependent variables are well formulated based on explanatory variables. Another class of statistical approaches consists of using simple parametric distributions, *e.g.*, the Weibull distribution for wind speed [97], or the beta distribution for solar irradiance [98] to model the density associated with the generative process. In this line, the (Gaussian) copula method has been widely used to model the spatial and temporal characteristics of wind [145] and PV generation [185]. For instance, the problem of generating probabilistic forecasts for the aggregated power of a set of renewable power plants harvesting different energy sources is addressed by Camal et al. [29].

Overall, these approaches usually make statistical assumptions increasing the difficulty to model the underlying stochastic process. The generated scenarios approximate the future uncertainty but cannot correctly describe all the salient features in the power output from renewable energy sources. *Deep learning* is one of the newest trends in artificial intelligence and machine learning to tackle the limitations of statistical methods with promising results across various application domains.

6.1.1 Related work

Recurrent neural networks (RNNs) are among the most famous deep learning techniques adopted in energy forecasting applications. A novel pooling-based deep recurrent neural network is proposed by Shi et al. [158] in the field of short-term household load forecasting. It outperforms statistical approaches such as autoregressive integrated moving average and classical RNN. A tailored forecasting tool, named encoder-decoder, is implemented in Dumas et al. [51] to compute intraday multi-output PV quantiles forecasts. Guidelines and best practices are developed by Hewamalage et al. [84] for forecasting practitioners on an extensive empirical study with an open-source software framework of existing RNN architectures. In the continuity, Toubreau et al. [166] implemented a bidirectional long short-term memory (BLSTM) architecture. It is trained using quantile regression and combined with a copula-based approach to generate scenarios. A scenario-based stochastic optimization case study compares this approach to other models regarding forecast quality and value. Finally, Salinas et al. [155] trained an autoregressive recurrent neural network on several real-world datasets. It produces accurate probabilistic forecasts with little or no hyper-parameter tuning.

Deep generative modeling is a class of techniques that trains deep neural networks to model the distribution of the observations. In recent years, there has been a growing interest in this field made possible by the appearance of large open-access datasets and breakthroughs in both general deep learning architectures and generative models. Several approaches exist such as energy-based models, variational autoencoders, generative adversarial networks, autoregressive models, normalizing flows, and numerous hybrid strategies. They all make trade-offs in terms of computation time, diversity, and architectural restrictions. We recommend two papers to get a broader knowledge of this field. (1) The comprehensive overview of generative modeling trends conducted by Bond-Taylor et al. [24]. It presents generative models to forecasting practitioners under a single cohesive statistical framework. (2) The thorough comparison of normalizing flows, variational autoencoders, and generative adversarial networks provided by Ruthotto and Haber [152]. It describes the advantages and disadvantages of each approach using numerical experiments in the field of computer vision. In the following, we focus on the applications of generative models in power systems.

In contrast to statistical approaches, deep generative models such as *Variational AutoEncoders* (VAEs) [102] and *Generative Adversarial Networks* (GANs) [74] directly learn a generative process of the data. They have demonstrated their effectiveness in many applications to compute accurate probabilistic forecasts, including power system applications. They both make probabilistic forecasts in the form of Monte Carlo

samples that can be used to compute consistent quantile estimates for all sub-ranges in the prediction horizon. Thus, they cannot suffer from the issue raised by Ordiano et al. [133] on the non-differentiable quantile loss function. Note: the generative models such as GANs and VAEs allow directly generating scenarios of the variable of interest. In contrast with methods that first compute weather scenarios to generate probabilistic forecasts such as implemented by Sun et al. [161] and Khoshrou and Pauwels [100]. A VAE composed of a succession of convolutional and feed-forward layers is proposed by Zhanga et al. [187] to capture the spatial-temporal complementary and fluctuant characteristics of wind and PV power with high model accuracy and low computational complexity. Both single and multi-output PV forecasts using a VAE are compared by Dairi et al. [37] to several deep learning methods such as LSTM, BLSTM, convolutional LSTM networks, and stacked autoencoders. The VAE consistently outperformed the other methods. A GAN is used by Chen et al. [31] to produce a set of wind power scenarios that represent possible future behaviors based only on historical observations and point forecasts. This method has a better performance compared to Gaussian Copula. A Bayesian GAN is introduced by Chen et al. [30] to generate wind and solar scenarios, and a progressive growing of GANs is designed by Yuan et al. [180] to propose a novel scenario forecasting method. In a different application, a GAN is implemented for building occupancy modeling without prior assumptions [32]. Finally, a conditional version of the GAN using several labels representing some characteristics of the demand is introduced by Lan et al. [110] to output power load data considering demand response programs.

Improved versions of GANs and VAEs have also been studied in the context of energy forecasting. The Wasserstein GAN enforces the Lipschitz continuity through a gradient penalty term (WGAN-GP), as the original GANs are challenging to train and suffer from mode collapse and over-fitting. Several studies applied this improved version in power systems: (1) a method using unsupervised labeling and conditional WGAN-GP models the uncertainties and variation in wind power [186]; (2) a WGAN-GP models both the uncertainties and the variations of the load [175]; (3) Jiang et al. [93] implemented scenario generation tasks both for a single site and for multiple correlated sites without any changes to the model structure. Concerning VAEs, they suffer from inherent shortcomings, such as the difficulties of tuning the hyper-parameters or generalizing a specific generative model structure to other databases. An improved VAE is proposed by Qi et al. [147] with the implementation of a β hyper-parameter into the VAE objective function to balance the two parts of the loss. This improved VAE is used to generate PV and power scenarios from historical values.

However, most of these studies did not benefit from conditional information such as weather forecasts to generate improved PV power, wind power, and load scenarios. In addition, to the best of our knowledge, only Ge et al. [63] compared NFs to these techniques for the generation of daily load profiles. Nevertheless, the comparison only considers quality metrics, and the models do not incorporate weather forecasts.

6.1.2 Research gaps and scientific contributions

This study investigates the implementation of *Normalizing Flows* [149, NFs] in power system applications. NFs define a new class of probabilistic generative models. They have gained increasing interest from the deep learning community in recent years. A NF learns a sequence of transformations, a *flow*, from a density known analytically, *e.g.*, a *Normal* distribution, to a complex target distribution. In contrast to other deep generative models, NFs can directly be trained by maximum likelihood estimation. They have proven to be an effective way to model complex data distributions with neural networks in many domains. First, speech synthesis [132]. Second, fundamental physics to increase the speed of gravitational wave inference by several orders of magnitude [78] or for sampling Boltzmann distributions of lattice field theories [4]. Finally, in the capacity firming framework by Dumas et al. [52].

This present work goes several steps further than Ge et al. [63] that demonstrated the competitiveness of NFs regarding GANs and VAEs for generating daily load profiles. First, we study the conditional version of these models to demonstrate that they can handle additional contextual information such as weather forecasts or geographical locations. Second, we extensively compare the model's performances both in terms of forecast value and quality. The forecast quality corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. The forecast value relates to the benefits of using forecasts in decision-making, such as participation in the electricity market. Third, we consider PV and wind generations in addition to load profiles. Finally, in contrast to the affine NFs used in their work, we rely on monotonic transformations, which are universal density approximators [88].

Given that Normalizing Flows are rarely used in the power systems community despite their potential, our main aim is to present this recent deep learning technique and demonstrate its interest and competitiveness with state-of-the-art generative models such as GANs and VAEs on a simple and easily reproducible case study. The research gaps motivating this study are three-fold:

1. To the best of our knowledge, only Ge et al. [63] compared NFs to GANs and VAEs for the generation of daily load profiles. Nevertheless, the comparison is only based on quality metrics, and the models do not take into account weather forecasts;
2. Most of the studies that propose or compare forecasting techniques only consider the forecast quality such as Ge et al. [63], Sun et al. [161], and Mashlakov et al. [121];
3. The conditional versions of the models are not always addressed, such as in Ge et al. [63]. However, weather forecasts are essential for computing accurate probabilistic forecasts.

With these research gaps in mind, the main contributions of this Chapter are three-fold:

1. We provide a fair comparison both in terms of quality and value with the state-of-the-art deep learning generative models, GANs and VAEs, using the open data of the Global Energy Forecasting Competition 2014 (GEFcom 2014) [86]. To the best of our knowledge, it is the first study that extensively compares the NFs, GANs, and VAEs on several datasets, PV generation, wind generation, and load with a proper assessment of the quality and value based on complementary metrics, and an easily reproducible case study;
2. We implement conditional generative models to compute improved weather-based PV, wind power, and load scenarios. In contrast to most of the previous studies that focused mainly on past observations;
3. Overall, we demonstrate that NFs are more accurate in quality and value, providing further evidence for deep learning practitioners to implement this approach in more advanced power system applications.

In addition to these contributions, this study also provides open-access to the Python code¹ to help the community to reproduce the experiments. Figure 6.2 provides the framework of the proposed method and Table 6.1 presents a comparison of the present study to three state-of-the-art papers using deep learning generative models to generate scenarios.

¹<https://github.com/jonathandumas/generative-models>

Criteria	[175]	[147]	[63]	study
GAN	✓	×	✓	✓
VAE	×	✓	✓	✓
NF	×	×	✓	✓
Number of models	4	1	3	3
PV	×	✓	×	✓
Wind power	×	✓	×	✓
Load	✓	~	✓	✓
Weather-based	✓	×	×	✓
Quality assessment	✓	✓	✓	✓
Quality metrics	5	3	5	8
Value assessment	×	✓	×	✓
Open dataset	~	×	✓	✓
Value replicability	-	~	-	✓
Open-access code	×	×	×	✓

Table 6.1 Comparison of the study’s contributions to three state-of-the-art studies using deep generative models.

✓: criteria fully satisfied, ~: criteria partially satisfied, ×: criteria not satisfied, ?: no information, -: not applicable. GAN: a GAN model is implemented; VAE: a VAE model is implemented; NF: a NF model is implemented; PV: PV scenarios are generated; Wind power: wind power scenarios are generated; Load: load scenarios are generated; Weather-based: the model generates weather-based scenarios; Quality assessment: a quality evaluation is conducted; Quality metrics: number of quality metrics considered; Value assessment: a value evaluation is considered with a case study; Open dataset: the data used for the quality and value evaluations are in open-access; Value replicability: the case study considered for the value evaluation is easily reproducible; Open-access code: the code used to conduct the experiments is in open-access. Note: the justifications are provided in Appendix 6.7.

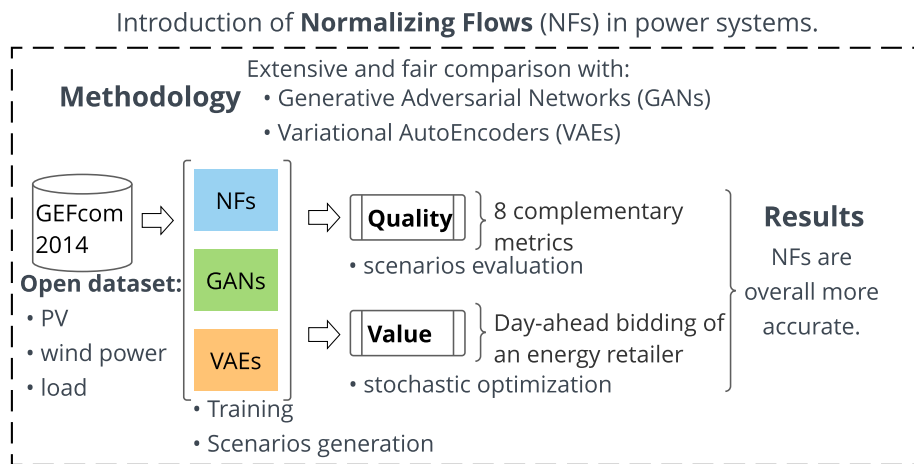


Fig. 6.2 The framework of the study.

The Chapter's primary purpose is to present and demonstrate the potential of NFs in power systems. A fair comparison is conducted both in terms of quality and value with the state-of-the-art deep learning generative models, GANs and VAEs, using the open data of the Global Energy Forecasting Competition 2014 [86]. The PV, wind power, and load datasets are used to assess the models. The quality evaluation is conducted by using eight complementary metrics, and the value assessment by considering the day-ahead bidding of an energy retailer using stochastic optimization. Overall, NFs tend to be more accurate both in terms of quality and value and are competitive with GANs and VAEs.

6.1.3 Applicability of the generative models

Probabilistic forecasting of PV, wind generation, electrical consumption, and electricity prices plays a vital role in renewable integration and power system operations. The deep learning generative models presented in this Chapter can be integrated into practical engineering applications. We present a non-exhaustive list of five applications in the following. (1) The forecasting module of an energy management system (EMS) [159]. Indeed, EMSs are used by several energy market players to operate various power systems such as a single renewable plant, a grid-connected or off-grid microgrid composed of several generations, consumption, and storage devices. An EMS is composed of several key modules: monitoring, forecasting, planning, control, *etc.* The forecasting module aims to provide the most accurate forecast of the variable of interest to be used as inputs of the planning and control modules. (2) Stochastic unit commitment models that employ scenarios to model the uncertainty of weather-dependent renewables. For instance, the optimal day-ahead scheduling and dispatch of a system composed of renewable plants, generators, and electrical demand are addressed by Camal et al. [29]. (3) Ancillary services market participation. A virtual power plant aggregating wind, PV, and small hydropower plants is studied by Camal et al. [29] to optimally bid on a day-ahead basis the energy and automatic frequency restoration reserve. (4) More generally, generative models can be used to compute scenarios for any variable of interest, *e.g.*, energy prices, renewable generation, loads, water inflow of hydro reservoirs, as long as data are available. (5) Finally, quantiles can be derived from scenarios and used in robust optimization models such as in the capacity firming framework [52].

6.1.4 Organization

The remainder of this Chapter is organized as follows. Section 6.2 presents the generative models implemented: NFs, GANs, and VAEs. Section 6.3 provides the quality and assessment methodology. Section 6.5 details empirical results on the GEFcom 2014 dataset, and Section 6.6 summarizes the main findings and highlights ideas for further work. Chapter 13 in Part II provides the forecast value assessment. Appendix 6.7 presents the justifications of Table 6.1, and Appendix 6.8 provides additional information on the generative models.

6.2 Background

This section formally introduces the conditional version of NFs, GANs, and VAEs implemented in this study. We assume the reader is familiar with the neural network's basics. However, for further information, Goodfellow et al. [76], Zhang et al. [184] provide a comprehensive introduction to modern deep learning approaches.

6.2.1 Multi-output forecasts using a generative model

Let us consider some dataset $\mathcal{D} = \{\mathbf{x}^i, \mathbf{c}^i\}_{i=1}^N$ of N independent and identically distributed samples from the joint distribution $p(\mathbf{x}, \mathbf{c})$ of two continuous variables X and C . X being the wind generation, PV generation, or load, and C the weather forecasts. They are both composed of T periods per day, with $\mathbf{x}^i := [x_1^i, \dots, x_T^i]^\top \in \mathbb{R}^T$ and $\mathbf{c}^i := [c_1^i, \dots, c_T^i]^\top \in \mathbb{R}^T$. The goal of this work is to generate multi-output weather-based scenarios $\hat{\mathbf{x}} \in \mathbb{R}^T$ that are distributed under $p(\mathbf{x}|\mathbf{c})$.

A generative model is a probabilistic model $p_\theta(\cdot)$, with parameters θ , that can be used as a generator of the data. Its purpose is to generate synthetic but realistic data $\hat{\mathbf{x}} \sim p_\theta(\mathbf{x}|\mathbf{c})$ whose distribution is as close as possible to the unknown data distribution $p(\mathbf{x}|\mathbf{c})$. In our application, it computes on a day-ahead basis a set of M scenarios at day $d - 1$ for day d

$$\hat{\mathbf{x}}_d^i := [\hat{x}_{d,1}^i, \dots, \hat{x}_{d,T}^i]^\top \in \mathbb{R}^T \quad i = 1, \dots, M. \quad (6.1)$$

For the sake of clarity, we omit the indexes d and i when referring to a scenario $\hat{\mathbf{x}}$ in the following.

6.2.2 Deep generative models

Figure 6.3 provides a high-level comparison of three categories of generative models considered in this Chapter: Normalizing Flows, Generative Adversarial Networks, and Variational AutoEncoders.

Normalizing flows

A normalizing flow is defined as a sequence of invertible transformations $f_k : \mathbb{R}^T \rightarrow \mathbb{R}^T$, $k = 1, \dots, K$, composed together to create an expressive invertible mapping $f_\theta := f_1 \circ \dots \circ f_K : \mathbb{R}^T \rightarrow \mathbb{R}^T$. This composed function can be used to perform density estimation, using f_θ to map a sample $\mathbf{x} \in \mathbb{R}^T$ onto a latent vector $\mathbf{z} \in \mathbb{R}^T$ equipped

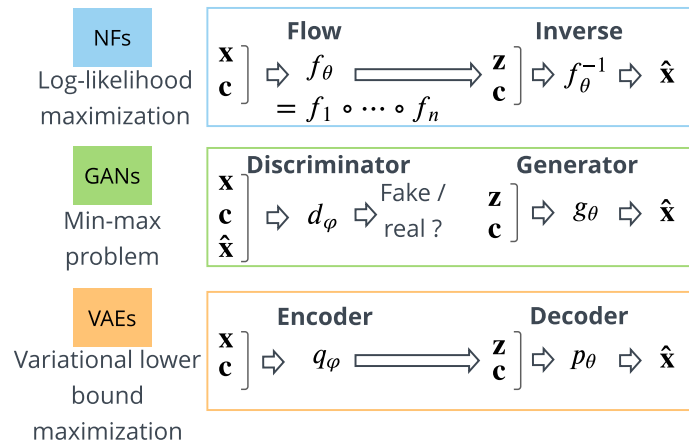


Fig. 6.3 High-level comparison of three categories of generative models considered in this Chapter: normalizing flows, generative adversarial networks, and variational autoencoders.

All models are conditional as they use the weather forecasts \mathbf{c} to generate scenarios $\hat{\mathbf{x}}$ of the distribution of interest \mathbf{x} : PV generation, wind power, load. Normalizing flows allow exact likelihood calculation. In contrast to generative adversarial networks and variational autoencoders, they explicitly learn the data distribution and provide direct access to the exact likelihood of the model's parameters. The inverse of the flow is used to generate scenarios. The training of generative adversarial networks relies on a min-max problem where the generator and the discriminator parameters are jointly optimized. The generator is used to compute the scenarios. Variational autoencoders indirectly optimize the log-likelihood of the data by maximizing the variational lower bound. The decoder computes the scenarios. Note: Section 6.2.3 provides a theoretical comparison of these models.

with a known and tractable probability density function p_z , *e.g.*, a Normal distribution. The transformation f_θ implicitly defines a density $p_\theta(\mathbf{x})$ that is given by the change of variables

$$p_\theta(\mathbf{x}) = p_z(f_\theta(\mathbf{x})) |\det J_{f_\theta}(\mathbf{x})|, \quad (6.2)$$

where J_{f_θ} is the Jacobian of f_θ regarding \mathbf{x} . The model is trained by maximizing the log-likelihood $\sum_{i=1}^N \log p_\theta(\mathbf{x}^i, \mathbf{c}^i)$ of the model's parameters θ given the dataset \mathcal{D} . For simplicity let us assume a single-step flow f_θ to drop the index k for the rest of the discussion.

In general, f_θ can take any form as long as it defines a bijection. However, a common solution to make the Jacobian computation tractable in (6.2) consists of implementing an *autoregressive* transformation [103], *i.e.*, such that f_θ can be rewritten as a vector of scalar bijections f^i

$$f_\theta(\mathbf{x}) := [f^1(x_1; h^1), \dots, f^T(x_T; h^T)]^\top, \quad (6.3a)$$

$$h^i := h^i(\mathbf{x}_{<i}; \varphi^i) \quad 2 \leq i \leq T, \quad (6.3b)$$

$$\mathbf{x}_{<i} := [x_1, \dots, x_{i-1}]^\top \quad 2 \leq i \leq T, \quad (6.3c)$$

$$h^1 \in \mathbb{R}, \quad (6.3d)$$

where $f^i(\cdot; h^i) : \mathbb{R} \rightarrow \mathbb{R}$ is partially parameterized by an autoregressive conditioner $h^i(\cdot; \varphi^i) : \mathbb{R}^{i-1} \rightarrow \mathbb{R}^{|h^i|}$ with parameters φ^i , and θ the union of all parameters φ^i .

There is a large choice of transformers f^i : affine, non-affine, integration-based, *etc.* This work implements an integration-based transformer by using the class of Unconstrained Monotonic Neural Networks (UMNN) [176]. It is a universal density approximator of continuous random variables when combined with autoregressive functions. The UMNN consists of a neural network architecture that enables learning arbitrary monotonic functions. It is achieved by parameterizing the bijection f^i as follows

$$f^i(x_i; h^i) = \int_0^{x_i} \tau^i(x_i, h^i) dt + \beta^i(h^i), \quad (6.4)$$

where $\tau^i(\cdot; h^i) : \mathbb{R}^{|h^i|+1} \rightarrow \mathbb{R}^+$ is the integrand neural network with a strictly positive scalar output, $h^i \in \mathbb{R}^{|h^i|}$ an embedding made by the conditioner, and $\beta^i(\cdot) : \mathbb{R}^{|h^i|} \rightarrow \mathbb{R}$ a neural network with a scalar output. The forward evaluation of f^i requires solving the integral (6.4) and is efficiently approximated numerically by using the Clenshaw-Curtis

quadrature. The pseudo-code of the forward and backward passes is provided by Wehenkel and Louppe [176].

Papamakarios et al. [136]’s Masked Autoregressive Network (MAF) is implemented to simultaneously parameterize the T autoregressive embeddings h^i of the flow (6.3). Then, the change of variables formula applied to the UMMN-MAF transformation results in the following log-density when considering weather forecasts

$$\log p_{\theta}(\mathbf{x}, \mathbf{c}) = \log p_z(f_{\theta}(\mathbf{x}, \mathbf{c})) |\det J_{f_{\theta}}(\mathbf{x}, \mathbf{c})|, \quad (6.5a)$$

$$= \log p_z(f_{\theta}(\mathbf{x}, \mathbf{c})) + \sum_{i=1}^T \log \tau^i(x_i, h^i(\mathbf{x}_{<i}), \mathbf{c}), \quad (6.5b)$$

that can be computed exactly and efficiently with a single forward pass. The UMMN-MAF approach implemented is referred to as NF in the rest of the Chapter. Figure 6.4 depicts the process of conditional normalizing flows with a three-step NF for PV generation. Note: Appendix 6.8.1 provides additional information on NFs.

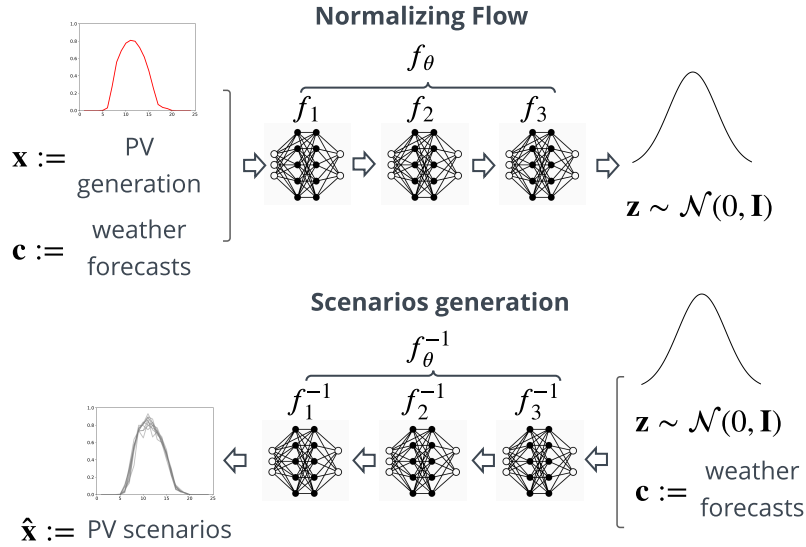


Fig. 6.4 The process of conditional normalizing flows is illustrated with a three-step NF for PV generation.

The model f_{θ} is trained by maximizing the log-likelihood of the model’s parameters θ given a dataset composed of PV observations and weather forecasts. Recall f_{θ} defines a bijection between the variable of interest \mathbf{x} , PV generation, and a Normal distribution \mathbf{z} . Then, the PV scenarios $\hat{\mathbf{x}}$ are generated by using the inverse of f_{θ} that takes as inputs samples from the Normal distribution \mathbf{z} and the weather forecasts \mathbf{c} .

Variational autoencoders

A VAE is a deep latent variable model composed of an *encoder* and a *decoder* which are jointly trained to maximize a lower bound on the likelihood. The encoder $q_\varphi(\cdot) : \mathbb{R}^T \times \mathbb{R}^{|\mathbf{c}|} \rightarrow \mathbb{R}^d$ approximates the intractable posterior $p(\mathbf{z}|\mathbf{x}, \mathbf{c})$, and the decoder $p_\theta(\cdot) : \mathbb{R}^d \times \mathbb{R}^{|\mathbf{c}|} \rightarrow \mathbb{R}^T$ the likelihood $p(\mathbf{x}|\mathbf{z}, \mathbf{c})$ with $\mathbf{z} \in \mathbb{R}^d$. Maximum likelihood is intractable as it would require marginalizing with respect to all possible realizations of the latent variables \mathbf{z} . Kingma and Welling [102] addressed this issue by maximizing the *variational lower bound* $\mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c})$ as follows

$$\log p_\theta(\mathbf{x}|\mathbf{c}) = KL[q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p(\mathbf{z}|\mathbf{x}, \mathbf{c})] + \mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c}), \quad (6.6a)$$

$$\geq \mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c}), \quad (6.6b)$$

$$\mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c}) := \mathbb{E}_{q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})} \left[\log \frac{p(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})}{q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})} \right], \quad (6.6c)$$

as the Kullback-Leibler (KL) divergence [141] is non-negative. Appendix 6.8.2 details how to compute the gradients of $\mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c})$, and its exact expression for the implemented VAE composed of fully connected neural networks for both the encoder and decoder. Figure 6.5 depicts the process of a conditional variational autoencoder for PV generation.

Generative adversarial networks

GANs are a class of deep generative models proposed by Goodfellow et al. [74] where the key idea is the adversarial training of two neural networks, the *generator* and the *discriminator*, during which the generator learns iteratively to produce realistic scenarios until they cannot be distinguished anymore by the discriminator from real data. The generator $g_\theta(\cdot) : \mathbb{R}^d \times \mathbb{R}^{|\mathbf{c}|} \rightarrow \mathbb{R}^T$ maps a latent vector $\mathbf{z} \in \mathbb{R}^d$ equipped with a known and tractable prior probability density function $p(\mathbf{z})$, *e.g.*, a Normal distribution, onto a sample $\mathbf{x} \in \mathbb{R}^T$, and is trained to fool the discriminator. The discriminator $d_\phi(\cdot) : \mathbb{R}^T \times \mathbb{R}^{|\mathbf{c}|} \rightarrow [0, 1]$ is a classifier trained to distinguish between true samples \mathbf{x} and generated samples $\hat{\mathbf{x}}$. Goodfellow et al. [74] demonstrated that solving the following min-max problem

$$\theta^* = \arg \min_{\theta} \max_{\phi} V(\phi, \theta), \quad (6.7)$$

where $V(\phi, \theta)$ is the value function, recovers the data generating distribution if $g_\theta(\cdot)$ and $d_\phi(\cdot)$ are given enough capacity. The state-of-the-art conditional Wasserstein GAN

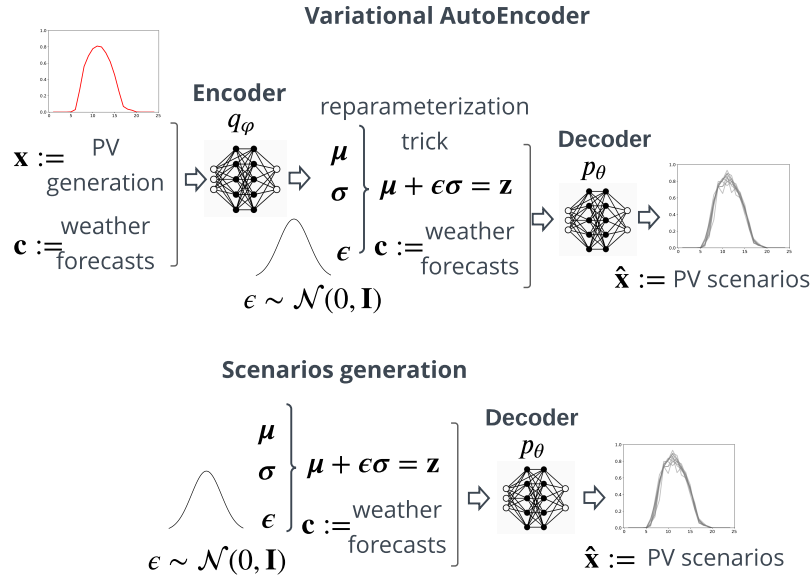


Fig. 6.5 The process of the conditional variational autoencoder is illustrated for PV generation.

The VAE is trained by maximizing the variational lower bound given a dataset composed of PV observations and weather forecasts. The encoder q_ϕ maps the variable of interest \mathbf{x} to a latent space \mathbf{z} . The decoder p_θ generates the PV scenarios $\hat{\mathbf{x}}$ by taking as inputs samples \mathbf{z} from the latent space and the weather forecasts \mathbf{c} .

with gradient penalty (WGAN-GP) proposed by Gulrajani et al. [79] is implemented with $V(\phi, \theta)$ defined as

$$V(\phi, \theta) = - \left(\mathbb{E}_{\tilde{\mathbf{x}}} [d_\phi(\hat{\mathbf{x}}|\mathbf{c})] - \mathbb{E}_{\mathbf{x}} [d_\phi(\mathbf{x}|\mathbf{c})] + \lambda \text{GP} \right), \quad (6.8a)$$

$$\text{GP} = \mathbb{E}_{\tilde{\mathbf{x}}} \left[\left(\|\nabla_{\tilde{\mathbf{x}}} d_\phi(\tilde{\mathbf{x}}|\mathbf{c})\|_2 - 1 \right)^2 \right], \quad (6.8b)$$

where $\tilde{\mathbf{x}}$ is implicitly defined by sampling convex combinations between the data and the generator distributions $\tilde{\mathbf{x}} = \rho \hat{\mathbf{x}} + (1 - \rho)\mathbf{x}$ with $\rho \sim \mathbb{U}(0, 1)$. The WGAN-GP constrains the gradient norm of the discriminator's output with respect to its input to enforce the 1-Lipschitz conditions. This strategy differs from the weight clipping of WGAN that sometimes generates only poor samples or fails to converge. Appendix 6.8.3 details the successive improvements from the original GAN to the WGAN and the final WGAN-GP implemented, referred to as GAN in the rest of the Chapter. Figure 6.6 depicts the process of a conditional generative adversarial network for PV generation.

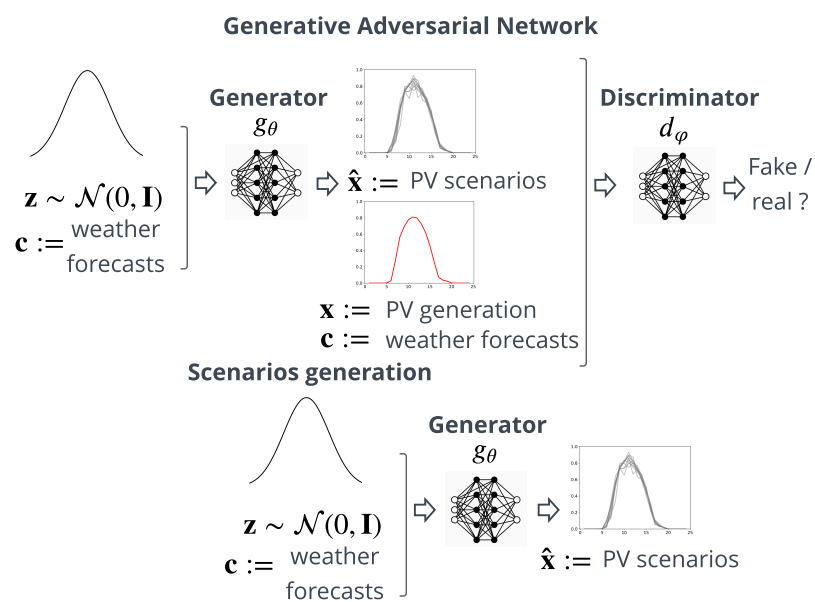


Fig. 6.6 The process of the conditional generative adversarial network is illustrated for PV generation.

The GAN is trained by solving a min-max problem given a dataset composed of PV observations \mathbf{x} and weather forecasts. The generator g_θ computes PV scenarios $\hat{\mathbf{x}}$ by taking as inputs samples from the Normal distribution \mathbf{z} and the weather forecasts \mathbf{c} , and the decoder d_ϕ tries to distinguish true data from scenarios.

6.2.3 Theoretical comparison

Normalizing flows are a generative model that allows exact likelihood calculation. They are efficiently parallelizable and offer a valuable latent space for downstream tasks. In contrast to GANs and VAEs, NFs explicitly learn the data distribution and provide direct access to the exact likelihood of the model's parameters, hence offering a sound and direct way to optimize the network parameters [177]. However, NFs suffer from drawbacks [24]. One disadvantage of requiring transformations to be invertible is that the input dimension must be equal to the output dimension, making the model difficult to train or inefficient. Each transformation must be sufficiently expressive while being easily invertible to compute the Jacobian determinant efficiently. The first issue is also raised by Ruthotto and Haber [152] where ensuring sufficient similarity of the distribution of interest and the latent distribution is of high importance to obtain meaningful and relevant samples. However, in our numerical simulations, we did not encounter this problem. Concerning the second issue, the UMNN-MAF transformation provides an expressive and effective way of computing the Jacobian.

VAEs indirectly optimize the log-likelihood of the data by maximizing the variational lower bound. The advantage of VAEs over NFs is their ability to handle non-invertible generators and the arbitrary dimension of the latent space. However, it has been observed that when applied to complex datasets such as natural images, VAEs samples tend to be unrealistic. There is evidence that the limited approximation of the true posterior, with a common choice being a normally distributed prior with diagonal covariance, is the root cause [190]. This statement comes from the field of computer vision. However, it may explain the shape of the scenarios observed in our numerical experiments in Section 6.5.

The training of GANs relies on a min-max problem where the generator and the discriminator are jointly optimized. Therefore, it does not rely on estimates of the likelihood or latent variable. The adversarial nature of GANs makes them notoriously difficult to train due to the saddle point problem [7]. Another drawback is the mode collapsing, where one network stays in bad local minima, and only a small subset of the data distribution is learned. Several improvements have been designed to address these issues, such as the Wasserstein GAN with gradient penalty. Thus, GANs models are widely used in computer vision and power systems. However, most GAN approaches require cumbersome hyperparameter tuning to achieve similar results to VAEs or NFs. In our numerical simulations, the GAN is highly sensitive to hyperparameter variations, which is consistent with [152].

Each method has its advantages and drawbacks and makes trade-offs in terms of computing time, hyper-parameter tuning, architecture complexity, *etc.* Therefore, the choice of a particular method is dependent on the user criteria and the dataset considered. In addition, the challenges of power systems are different from computer vision. Therefore, the limitations established in the computer vision literature such as Bond-Taylor et al. [24] and Ruthotto and Haber [152] must be addressed with caution. Therefore, we encourage the energy forecasting practitioners to test and compare these methods in power systems applications.

6.3 Quality assessment

For predictions in any form, one must differentiate between their quality and their value [124]. Forecast quality corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. Forecast value relates, instead, to the benefits from using forecasts in a decision-making process, such as participation in the electricity market. The forecast value is assessed in Part II Chapter 13 by considering the day-ahead market scheduling of electricity aggregators, such as energy retailers or generation companies.

Evaluating and comparing generative models remains a challenging task. Several measures have been introduced with the emergence of new models, particularly in the field of computer vision. However, there is no consensus or guidelines as to which metric best captures the strengths and limitations of models. Generative models need to be evaluated directly to the application they are intended for [165]. Indeed, good performance to one criterion does not imply good performance to the other criteria. Several studies propose metrics and make attempts to determine the pros and cons. We selected two that provide helpful information. (1) 24 quantitative and five qualitative measures for evaluating generative models are reviewed and compared by Borji [26] with a particular emphasis on GAN-derived models. (2) several representative sample-based evaluation metrics for GANs are investigated by Huang et al. [89] where the kernel Maximum Mean Discrepancy (MMD) and the 1-Nearest-Neighbour (1-NN) two-sample test seem to satisfy most of the desirable properties. The key message is to combine several complementary metrics to assess the generative models. Some of the metrics proposed are related to image generation and cannot directly be transposed to energy forecasting.

Therefore, we used eight complementary quality metrics to conduct a relevant quality analysis inspired by the energy forecasting and computer vision fields. They

can be divided into four groups: (1) the *univariate* metrics composed of the continuous ranked probability score, the quantile score, and the reliability diagram. They can only assess the quality of the scenarios with respect to their marginals; (2) the *multivariate* metrics are composed of the energy and the variogram scores. They can directly assess multivariate scenarios; (3) the *specific* metrics composed of a classifier-based metric and the correlation matrix between scenarios for a given context; (4) the Diebold and Mariano statistical test. The univariate and multivariate metrics are already defined in Chapter 3. Therefore, only the specific metrics and statistical test are introduced in this Section.

6.3.1 Classifier-based metric

Modern binary classifiers can be easily turned into robust two-sample tests where the goal is to assess whether two samples are drawn from the same distribution [113]. In other words, it aims at assessing whether a generated scenario can be distinguished from an observation. To this end, the generator is evaluated on a held-out testing set that is split into a testing-train and testing-test subsets. The testing-train set is used to train a classifier, which tries to distinguish generated scenarios from the actual distribution. Then, the final score is computed as the performance of this classifier on the testing-test set.

In principle, any binary classifier can be adopted for computing classifier two-sample tests (C2ST). A variation of this evaluation methodology is proposed by Huang et al. [89] and is known as the 1-Nearest Neighbor (NN) classifier. The advantage of using 1-NN over other classifiers is that it requires no special training and little hyper-parameter tuning. This process is conducted as follows. Given two sets of observations S_r and generated S_g samples with the same size, *i.e.*, $|S_r| = |S_g|$, it is possible to compute the leave-one-out (LOO) accuracy of a 1-NN classifier trained on S_r and S_g with positive labels for S_r and negative labels for S_g . The LOO accuracy can vary from 0 % to 100 %. The 1-NN classifier should yield a 50 % LOO accuracy when $|S_r| = |S_g|$ is large. This is achieved when the two distributions match. Indeed, the level 50 % happens when a label is randomly assigned to a generated scenario. It means the classifier is not capable of discriminating generated scenarios from observations. If the generative model over-fits S_g to S_r , *i.e.*, $S_g = S_r$, and the accuracy would be 0 %. On the contrary, if it generates widely different samples than observations, the performance should be 100 %. Therefore, the closer the LOO accuracy is to 1, the higher the degree of under-fitting of the model. The closer the LOO accuracy is to 0, the higher the degree of over-fitting of the model. The C2ST approach using LOO

with 1-NN is adopted by Qi et al. [147] to assess the PV and wind power scenarios of a β VAE.

However, this approach has several limitations. First, it uses the testing set to train the classifier during the LOO. Second, the 1-NN is very sensitive to outliers as it simply chose the closest neighbor based on distance criteria. This behavior is amplified when combined with the LOO where the testing-test set is composed of only one sample. Third, the euclidian distance cannot deal with a context such as weather forecasts. Therefore, we cannot use a conditional version of the 1-NN using weather forecasts to classify weather-based renewable generation and the observations. Fourth, C2ST with LOO cannot provide ROC curve but only accuracy scores. An essential point about ROC graphs is that they measure the ability of a classifier to produce good relative instance scores. In our case, we are interested in discriminating the generated scenarios from the observations, and the ROC provides more information than the accuracy metric to achieve this goal. A standard method to reduce ROC performance to a single scalar value representing expected performance is to calculate the area under the ROC curve, abbreviated AUC. The AUC has an important statistical property: it is equivalent to the probability that the classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance [58].

To deal with these issues, we decided to modify this classifier-based evaluation by conducting the C2ST as follows: (1) the scenarios generated on the learning set are used to train the classifier using the C2ST. Therefore, the classifier uses the entire testing set and can compute ROC; (2) the classifier is an Extra-Trees classifier that can deal with context such as weather forecasts.

More formally, for a given generative model g , the following steps are conducted:

1. Initialization step: the generative model g has been trained on the LS and has generated M weather-based scenarios per day of both the LS and TS: $\{\hat{\mathbf{x}}_{\text{LS}}^i\}_{i=1}^M := \cup_{d \in \text{LS}} \{\hat{\mathbf{x}}_d^i\}_{i=1}^M$ and $\{\hat{\mathbf{x}}_{\text{TS}}^i\}_{i=1}^M := \cup_{d \in \text{TS}} \{\hat{\mathbf{x}}_d^i\}_{i=1}^M$. For the sake of clarity the index g is omitted, but both of these sets are dependent on model g .
2. M pairs of learning and testing sets are built with an equal proportion of generated scenarios and observations: $\mathcal{D}_{\text{LS}}^i := \left\{ \{\hat{\mathbf{x}}_{\text{LS}}^i, 0\} \cup \{\{\mathbf{x}_{\text{LS}}^i, 1\}\} \right\}$ and $\mathcal{D}_{\text{TS}}^i = \left\{ \{\hat{\mathbf{x}}_{\text{TS}}^i, 0\} \cup \{\{\mathbf{x}_{\text{TS}}^i, 1\}\} \right\}$. Note: $|\mathcal{D}_{\text{LS}}^i| = 2|\text{LS}|$ and $|\mathcal{D}_{\text{TS}}^i| = 2|\text{TS}|$.
3. For each pair of learning and testing sets $\{\mathcal{D}_{\text{LS}}^i, \mathcal{D}_{\text{TS}}^i\}_{i=1}^M$ a classifier d_g^i is trained and makes predictions.

4. The ROC_g^i curves and corresponding AUC_g^i are computed for $i = 1, \dots, M$.

This classifier-based methodology is conducted for all models g , and the results are compared. Figure 6.7 depicts the overall approach. The classifiers d_g^i are all Extra-Trees classifier made of 1000 unconstrained trees with the hyper-parameters "max_depth" set to "None", and "n_estimators" to 1 000.

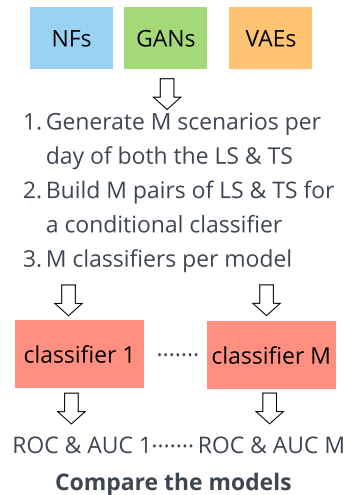


Fig. 6.7 Classifier-based metric methodology.

Each generative model generates M scenarios per day of the learning and testing sets. They are used to build M pairs of learning and testing sets for a conditional classifier by including an equal proportion of observations and weather forecasts. M conditional classifiers, per model, are trained and make predictions. The M ROC and AUC are computed per model, and the results are compared.

6.3.2 Correlation matrix between scenarios

The second specific metric consists of computing the correlation matrix between the scenarios generated for given weather forecasts. Formally, let $\{\hat{\mathbf{x}}^i\}_{i=1}^M$ be the set of M scenarios generated for a given day of the testing set. It is a matrix $(M \times 24)$ where each row is a scenario. Then, the Pearson's correlation coefficients are computed into a correlation matrix (24×24) . This metric indicates the variety of scenario shapes.

6.3.3 Diebold-Mariano test

Using relevant metrics to assess the forecast quality is essential. However, it is also necessary to analyze whether any difference in accuracy is statistically significant.

Indeed, when different models have almost identical values in the selected error measures, it is not easy to draw statistically significant conclusions on the outperformance of the forecasts of one model by those of another. The Diebold-Mariano (DM) test [43] is probably the most commonly used statistical testing tool to evaluate the significance of differences in forecasting accuracy. It is model-free, *i.e.*, it compares the forecasts of models, and not models themselves. The DM test is used in this study to assess the CRPS, QS, ES, and VS metrics. The CRPS and QS are univariate scores, and a value of CRPS and QS is computed per marginal (time period of the day). Therefore, the multivariate variant of the DM test is implemented following Ziel and Weron [192], where only one statistic for each pair of models is computed based on the 24-dimensional vector of errors for each day.

For a given day d of the testing set, let $\epsilon_d \in \mathbb{R}$ be the error computed by an arbitrary forecast loss function of the observation and scenarios. The test consists of computing the difference between the errors of the pair of models g and h over the testing set

$$\Delta(g, h)_d = \epsilon_d^g - \epsilon_d^h, \quad \forall d \in \text{TS}, \quad (6.9)$$

and to perform an asymptotic z -test for the null hypothesis that the expected forecast error is equal and the mean of differential loss series is zero $\mathbb{E}[\Delta(g, h)_d] = 0$. It means there is no statistically significant difference in the accuracy of the two competing forecasts. The statistic of the test is deduced from the asymptotically standard normal distribution as follows

$$\text{DM}(g, h) = \sqrt{\#\text{TS}} \frac{\hat{\mu}}{\hat{\sigma}}, \quad (6.10)$$

with $\#\text{TS}$ the number of days of the testing set, $\hat{\mu}$ and $\hat{\sigma}$ the sample mean and the standard deviation of $\Delta(g, h)$. Under the assumption of covariance stationarity of the loss differential series $\Delta(g, h)_d$, the DM statistic is asymptotically standard normal. The lower the p -value, *i.e.*, the closer it is to zero, the more the observed data is inconsistent with the null hypothesis: $\mathbb{E}[\Delta(g, h)_d] < 0$ the forecasts of the model h are more accurate than those of model g . If the p -value is less than the commonly accepted level of 5 %, the null hypothesis is typically rejected. It means that the forecasts of model g are significantly more accurate than those of model h .

When considering the ES or VS scores, there is a value per day of the testing set ES_d or VS_d . In this case, $\epsilon_d = \text{ES}_d$ or $\epsilon_d = \text{VS}_d$. However, when considering the CRPS or QS, there is a value per marginal and per day of the testing set $\text{CRPS}_{d,k}$ or $\text{QS}_{d,k}$. A solution consists of computing 24 independent tests, one for each hour of the day. Then, to compare the models based on the number of hours for which the

predictions of one model are significantly better than those of another. Another way consists of a multivariate variant of the DM-test with the test performed jointly for all hours using the multivariate loss differential series. In this case, for a given day d , $\epsilon_d^g = [\epsilon_{d,1}^g, \dots, \epsilon_{d,24}^g]^\top$, $\epsilon_d^h = [\epsilon_{d,1}^h, \dots, \epsilon_{d,24}^h]^\top$ are the vectors of errors for a given metric of models g and h , respectively. Then the multivariate loss differential series

$$\Delta(g, h)_d = \|\epsilon_d^g\|_1 - \|\epsilon_d^h\|_1, \quad (6.11)$$

defines the differences of errors using the $\|\cdot\|_1$ norm. Then, the p -value of two-sided DM tests is computed for each model pair as described above. The univariate version of the test has the advantage of providing a more profound analysis as it indicates which forecast is significantly better for which hour of the day. The multivariate version enables a better representation of the results as it summarizes the comparison in a single p -value, which can be conveniently visualized using heat maps arranged as chessboards. In this study, we decided to adopt the multivariate DM-test for the CRPS and QS.

6.4 Case study

The quality and value evaluations of the models are conducted on the load, wind, and PV tracks of the open-access GEFCom 2014 dataset [86], composed of one, ten, and three zones, respectively. See Chapter 13 for the value evaluation and the energy retailer problem statement. Figure 6.8 depicts the methodology to assess both the quality and value of the GAN, VAE and NF models implemented in this study.

6.4.1 Implementation details

By appropriate normalization, we standardize the weather forecasts to have a zero mean and unit variance. Table 6.2 provides a summary of the implementation details described in what follows. For the sake of proper model training and evaluation, the dataset is divided into three parts per track considered: learning, validation, and testing sets. The learning set (LS) is used to fit the models, the validation set (VS) to select the optimal hyper-parameters, and the testing set (TS) to assess the forecast quality and value. The number of samples ($\#$), expressed in days, of the VS and TS, is $50 \cdot n_z$, with n_z the number of zones of the track considered. The 50 days are selected randomly from the dataset, and the learning set is composed of the remaining part with $D \cdot n_z$ samples, where D is provided for each track in Table 6.2. The NF, VAE,

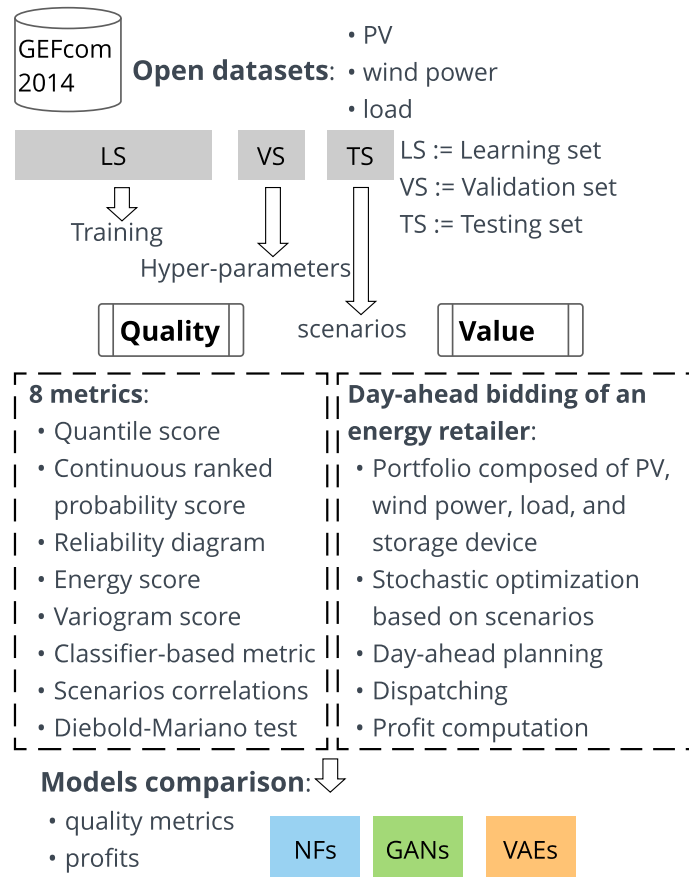


Fig. 6.8 Methodology to assess the quality and value of the GAN, VAE, and NF models implemented in this study.

The PV, wind power, and load datasets of the open-access Global Energy Forecasting Competition 2014 are divided into three parts: learning, validation, and testing sets. The learning set is used to train the models, the validation set to select the optimal hyper-parameters, and the testing set to conduct the numerical experiments. The quality and value of the models are assessed by using the scenarios generated on the testing set. The quality evaluation consists of eight complementary metrics. The value assessment uses the simple and easily reproducible case study of the day-ahead bidding of an energy retailer. The energy retailer portfolio is composed of PV, wind power generation, load, and a storage system device. The retailer bids on the day-ahead market by computing a planning based on stochastic optimization. The dispatch is computed by using the observations of the PV generation, wind power, and load. Then, the profits are evaluated and compared.

and GAN use the weather forecasts as inputs to generate on a day-ahead basis M scenarios $\hat{\mathbf{x}} \in \mathbb{R}^T$.

Wind track

The zonal \mathbf{u}^{10} , \mathbf{u}^{100} and meridional \mathbf{v}^{10} , \mathbf{v}^{100} wind components at 10 and 100 meters are selected, and six features are derived following the formulas provided by Landry et al. [111] to compute the wind speed \mathbf{ws}^{10} , \mathbf{ws}^{100} , energy \mathbf{we}^{10} , \mathbf{we}^{100} and direction \mathbf{wd}^{10} , \mathbf{wd}^{100} at 10 and 100 meters

$$\mathbf{ws} = \sqrt{\mathbf{u} + \mathbf{v}}, \quad (6.12a)$$

$$\mathbf{we} = \frac{1}{2} \mathbf{ws}^3, \quad (6.12b)$$

$$\mathbf{wd} = \frac{180}{\pi} \arctan(\mathbf{u}, \mathbf{v}). \quad (6.12c)$$

For each generative model, the wind zone is taken into account with one hot-encoding variable Z_1, \dots, Z_{10} , and the wind feature input vector for a given day d is

$$\mathbf{c}_d^{\text{wind}} = [\mathbf{u}_d^{10}, \mathbf{u}_d^{100}, \mathbf{v}_d^{10}, \mathbf{v}_d^{100}, \mathbf{ws}_d^{10}, \mathbf{ws}_d^{100}, \mathbf{we}_d^{10}, \mathbf{we}_d^{100}, \mathbf{wd}_d^{10}, \mathbf{wd}_d^{100}, Z_1, \dots, Z_{10}], \quad (6.13)$$

of dimension $n_f \cdot T + n_z = 10 \cdot 24 + 10$.

PV track

The solar irradiation \mathbf{I} , the air temperature \mathbf{T} , and the relative humidity \mathbf{rh} are selected, and two features are derived by computing \mathbf{I}^2 and \mathbf{IT} . For each generative model, the PV zone is taken into account with one hot-encoding variable Z_1, Z_2, Z_3 , and the PV feature input vector for a given day d is

$$\mathbf{c}_d^{\text{PV}} = [\mathbf{I}_d, \mathbf{T}_d, \mathbf{rh}_d, \mathbf{I}_d^2, \mathbf{IT}_d, Z_1, Z_2, Z_3], \quad (6.14)$$

of dimension $n_f \cdot T + n_z$. For practical reasons, the periods where the PV generation is always 0, across all zones and days, are removed, and the final dimension of the input feature vector is $n_f \cdot T + n_z = 5 \cdot 16 + 3$.

Load track

The 25 weather station temperature $\mathbf{w}_1, \dots, \mathbf{w}_{25}$ forecasts are used. There is only one zone, and the load feature input vector for a given day d is

$$\mathbf{c}_d^{\text{load}} = [\mathbf{w}_1, \dots, \mathbf{w}_{25}], \quad (6.15)$$

of dimension $n_f \cdot T = 25 \cdot 24$.

	Wind	PV	Load
T periods	24	16	24
n_z zones	10	3	—
n_f features	10	5	25
\mathbf{c}_d dimension	$n_f \cdot T + n_z$	$n_f \cdot T + n_z$	$n_f \cdot T$
# LS (days)	$631 \cdot n_z$	$720 \cdot n_z$	1999
# VS/TS (days)	$50 \cdot n_z$	$50 \cdot n_z$	50

Table 6.2 Dataset and implementation details.

Each dataset is divided into three parts: learning, validation, and testing sets. The number of samples (#) is expressed in days and is set to 50 days for the validation and testing sets. T is the number of periods per day considered, n_z the number of zones of the dataset, n_f the number of weather variables used, and \mathbf{c}_d is the dimension of the conditional vector for a given day that includes the weather forecasts and the one hot-encoding variables when there are several zones. Note: the days of the learning, validation, and testing sets are selected randomly.

6.4.2 Hyper-parameters

Table 6.3 provides the hyper-parameters of the NF, VAE, and GAN implemented. The Adam optimizer [101] is used to train the generative models with a batch size of 10% of the learning set. The NF implemented is a one-step monotonic normalizer using the UMNN-MAF². The embedding size $|h^i|$ is set to 40, and the embedding neural network is composed of l layers of n neurons ($l \times n$). The same integrand neural network $\tau^i(\cdot) \forall i = 1, \dots, T$ is used and composed of 3 layers of $|h^i|$ neurons (3×40). Both the encoder and decoder of the VAE are feed-forward neural networks ($l \times n$), ReLU activation functions for the hidden layers, and no activation function for

²<https://github.com/AWehenkel/Normalizing-Flows>

	Wind	PV	Load
Embedding Net	4×300	4×300	4×300
Embedding size	40	40	40
(a)Integrand Net	3×40	3×40	3×40
Weight decay	$5 \cdot 10^{-4}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-4}$
Learning rate	10^{-4}	$5 \cdot 10^{-4}$	10^{-4}
Latent dimension	20	40	5
(b)E/D Net	1×200	2×200	1×500
Weight decay	$10^{-3.4}$	$10^{-3.5}$	10^{-4}
Learning rate	$10^{-3.4}$	$10^{-3.3}$	$10^{-3.9}$
Latent dimension	64	64	256
(c)G/D Net	2×256	3×256	2×1024
Weight decay	10^{-4}	10^{-4}	10^{-4}
Learning rate	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$

Table 6.3 (a) NF, (b) VAE, and (c) GAN hyper-parameters.

The hyper-parameters selection is performed on the validation set using the Python library Weights & Biases [21]. This library is an experiment tracking tool for machine learning, making it easier to track experiments. The GAN model was the most time-consuming during this process, followed by the VAE and NF. Indeed, the GAN is highly sensitive to hyper-parameter modifications making it challenging to identify a relevant set of values. In contrast, the NF achieved satisfactory results, both in terms of scenarios shapes and quality, by testing only a few sets of hyper-parameter values.

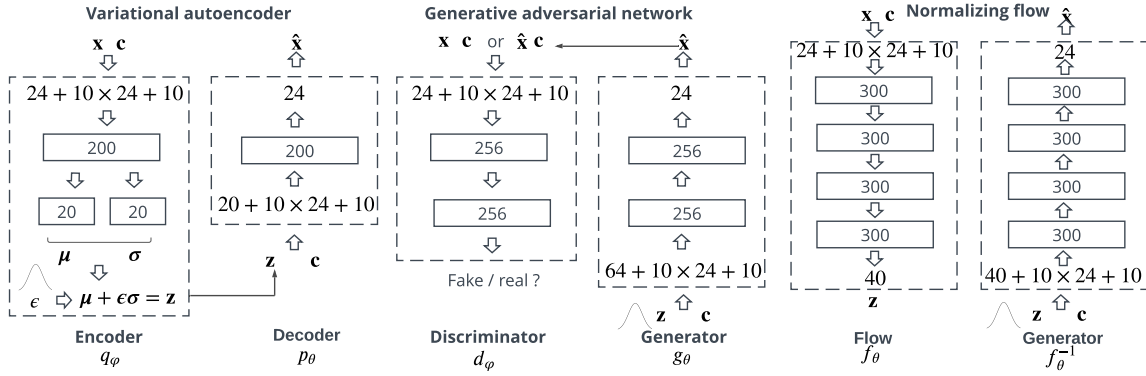


Fig. 6.9 Variational autoencoder, generative adversarial network, and normalizing flow structures implemented for the wind dataset.

VAE: both the encoder and decoder are feed-forward neural networks composed of one hidden layer with 200 neurons. Increasing the number of layers did not improve the results for this dataset. The latent space dimension is 20.

GAN: both the discriminator and generator are feed-forward neural networks composed of two hidden layers with 256 neurons. The latent space dimension is 64.

NF: a single-step monotonic normalizing flow is implemented with a feed-forward neural network composed of four hidden layers with 300 neurons. The latent space dimension is 40. Note: For clarity, the integrand network is not included but is a feed-forward neural network composed of three hidden layers with 40 neurons. Increasing the number of steps of the normalizing flow did not improve the results. The monotonic transformation is complex enough to capture the stochasticity of the variable of interest. However, when considering affine autoregressive normalizing flows, the number of steps should be generally higher. Numerical experiments indicated that a five-step autoregressive flow was required to achieve similar results for this dataset. Note: the results are not reported in this study for the sake of clarity

the output layer. Both the generator and discriminator of the GAN are feed-forward neural networks ($l \times n$). The activation functions of the hidden layers of the generator (discriminator) are ReLU (Leaky ReLU). The activation function of the discriminator output layer is ReLU, and there is no activation function for the generator output layer. The generator is trained once after the discriminator is trained five times to stabilize the training process, and the gradient penalty coefficient λ in (6.7) is set to 10 as suggested by Gulrajani et al. [79].

Figure 6.9 illustrates the VAE, GAN, and NF structures implemented for the wind dataset where the number of weather variables selected and the number of zones is 10, and 10, respectively. Recall, $\mathbf{c} :=$ weather forecasts, $\hat{\mathbf{x}} :=$ scenarios $\mathbf{x} :=$ wind power observations, $\mathbf{z} :=$ latent space variable, $\epsilon :=$ Normal variable (only for the VAE).

6.5 Quality results

First, a thorough comparison of the models is conducted on the wind track. Second, all tracks are considered for the sake of clarity. Note that the model ranking slightly differs depending on the track.

6.5.1 Wind track

In addition to the generative models, a naive approach is designed (RAND). The scenarios of the learning, validation and testing sets are sampled randomly from the learning, validation, and testing sets, respectively. Intuitively, it boils down to assume that past observations are repeated. These scenarios are realistic but may not be compatible with the context. Each model generates a set of 100 scenarios for each day of the testing set. And the metrics are computed. Figure 6.10 compares the QS, reliability diagram, and CRPS of the wind (markers), PV (plain), and load (dashed) tracks. Overall, for the wind track in terms of CRPS, QS, and reliability diagrams, the VAE achieves slightly better scores, followed by the NF and the GAN. The ES and VS multivariate scores confirm this trend with 54.82 and 18.87 for the VAE *vs* 56.71 and 18.54 for the NF, respectively.

Figure 6.13 provides the results of the DM tests for these metrics. The heat map indicates the range of the p -values. The closer they are to zero, dark green, the more significant the difference between the scores of two models for a given metric. The statistical threshold is five %, but the scale color is capped at ten % for a better exposition of the relevant results. For instance, when considering the DM test for the RAND CRPS, all the columns of the RAND row are in dark green, indicating that the RAND scenarios are always significantly outperformed by the other models. These DM tests confirm that the VAE outperforms the NF for the wind track considering these metrics. Then, the NF is only outperformed by the VAE and the GAN by both the VAE and NF. These results are consistent with the classifier-based metric depicted in Figure 6.11a, where the VAE is the best to mislead the classifier followed by the NF, and GAN.

The left part of Figure 6.12 provides 50 scenarios, (a) NF, (c) GAN and (e) VAE, generated for a given day selected randomly from the testing set. Notice how the shape of the NF's scenarios differs significantly from the GAN and VAE as they tend to be more variable with no identifiable trend. In contrast, the VAE and GAN scenarios seem to differ mainly in nominal power but have similar shapes. This behavior is even more pronounced for the GAN, where the scenarios rarely crossed over the periods.

For instance, there is a gap in generation around periods 17 and 18 where all the GAN's scenarios follow this trend. These observations are confirmed by computing the corresponding time correlation matrices, depicted by the right part of Figure 6.12 demonstrating there is no correlation between NF's scenarios. On the contrary, the VAE and GAN correlation matrices tend to be similar with a time correlation of the scenarios over a few periods, with more correlated periods when considering the GAN. This difference in the scenario's shape is striking and not necessarily captured by metrics such as the CRPS, QS, or even the classifier-based metric and is also observed on the PV and load tracks, as explained in the next paragraph.

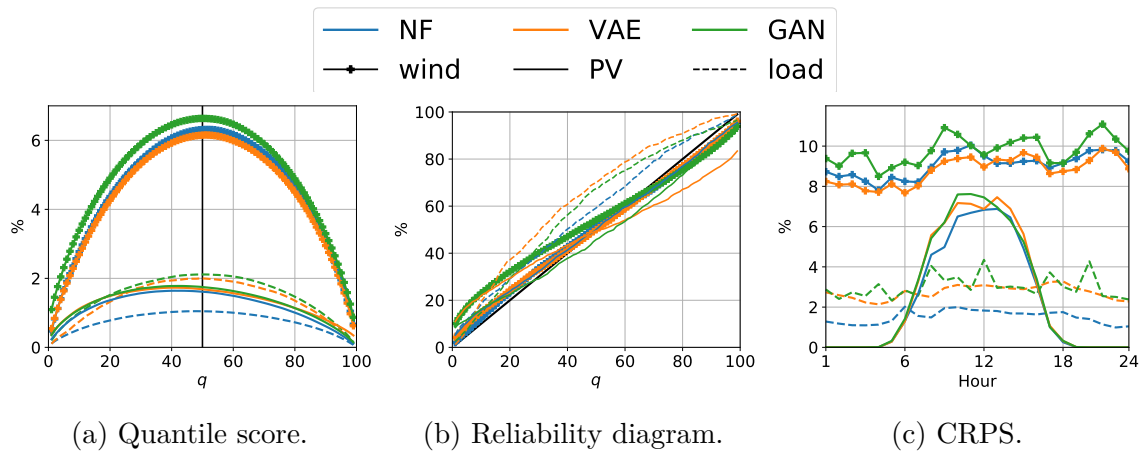


Fig. 6.10 Quality standard metrics comparison on the wind (markers), PV (plain), and load (dashed) tracks.

Quantile score (a): the lower and the more symmetrical, the better. Note: the quantile score has been averaged over the marginals (the 24 time periods of the day). Reliability diagram (b): the closer to the diagonal, the better. Continuous ranked probability score per marginal (c): the lower, the better. NF outperforms the VAE and GAN for both the PV and load tracks and is slightly outperformed by the VAE on the wind track. Note: all models tend to have more difficulties forecasting the wind power that seems less predictable than the PV generation or the load.

6.5.2 All tracks

Table 6.4 provides the averaged quality scores. The CRPS is averaged over the 24 time periods $\overline{\text{CRPS}}$. The QS over the 99 percentiles $\overline{\text{QS}}$. The MAE-r is the mean absolute error between the reliability curve and the diagonal, and $\overline{\text{AUC}}$ is the mean of the 50 AUC. Overall, for the PV and load tracks in CRPS, QS, reliability diagrams, AUC, ES, and VS, the NF outperforms the VAE and GAN and is slightly outperformed by the VAE on the wind track. On the load track, the VAE outperforms the GAN.

		NF	VAE	GAN	RAND
Wind	$\overline{\text{CRPS}}$	9.07	8.80	9.79	16.92
	$\overline{\text{QS}}$	4.58	4.45	4.95	8.55
	MAE-r	2.83	2.67	6.82	1.01
	$\overline{\text{AUC}}$	0.935	0.877	0.972	0.918
	ES	56.71	54.82	60.52	96.15
	VS	18.54	17.87	19.87	23.21
PV	$\overline{\text{CRPS}}$	2.35	2.60	2.61	4.92
	$\overline{\text{QS}}$	1.19	1.31	1.32	2.48
	MAE-r	2.66	9.04	4.94	3.94
	$\overline{\text{AUC}}$	0.950	0.969	0.997	0.947
	ES	23.08	24.65	24.15	41.53
	VS	4.68	5.02	4.88	13.40
Load	$\overline{\text{CRPS}}$	1.51	2.74	3.01	6.74
	$\overline{\text{QS}}$	0.76	1.39	1.52	3.40
	MAE-r	7.70	13.97	9.99	0.88
	$\overline{\text{AUC}}$	0.823	0.847	0.999	0.944
	ES	9.17	15.11	17.96	38.08
	VS	1.63	1.66	3.81	7.28

Table 6.4 Averaged quality scores per dataset.

The best performing deep learning generative model for each track is written in bold. The CRPS, QS, MAE-r, and ES are expressed in %. Overall, for both the PV and load tracks, the NF outperforms the VAE and GAN and slightly outperforms the VAE on the wind track.




Fig. 6.11 Wind, PV, and load tracks classifier-based metric.

Wind: the VAE (orange) is the best to mislead the classifier, followed by the NF (blue) and GAN (green). PV and load: the NF (blue) is the best to fake the classifier, followed by the VAE (orange) and the GAN (green). Note: there are 50 ROC curves depicted for each model, each corresponding to a scenario generated used as input of the classifier. It allows taking into account the variability of the scenarios to avoid having results dependent on a particular scenario.

However, the VAE and GAN achieved similar results on the PV track, and the GAN performed better in terms of ES and VS. These results are confirmed by the DM tests depicted in Figure 6.13. The classifier-based metric results for both the load and PV tracks, provided by Figure 6.11, confirm this trend where the NF is the best to trick the classifier followed by the VAE and GAN.

Similar to the wind track, the shape of the scenarios differs significantly between the NF and the other models for both the load and PV tracks as indicated by Figure 6.14. Note: the load track scenarios are highly correlated for both the VAE and GAN. Finally, Figure 6.15 provides the average of the correlation matrices over all days of the testing set for each dataset. The trend depicted above is confirmed. This difference between the NF and the other generative model may be explicated by the design of the methods. The NF explicitly learns the probability density function (PDF) of the multi-dimensional random variable considered. Thus, the NF scenarios are generated according to the learned PDF producing multiple shapes of scenarios. In contrast, the generator of the GAN is trained to fool the discriminator, and it may find a shape particularly efficient leading to a set of similar scenarios. Concerning the VAE, it is less obvious. However, by design, the decoder is trained to generate scenarios from the latent space assumed to follow a Gaussian distribution that may lead to less variability.

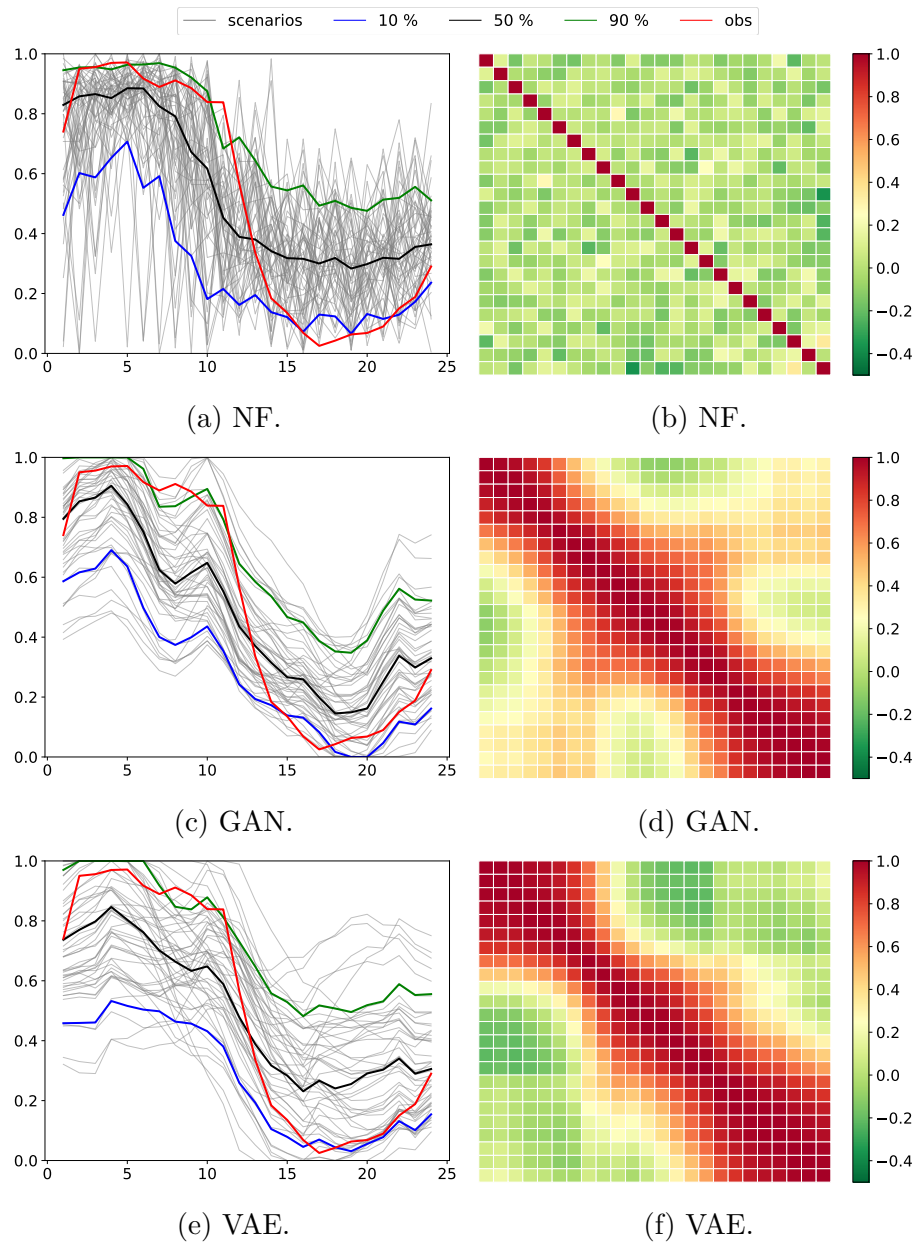


Fig. 6.12 Wind power scenarios shape comparison and analysis.

Left part (a) NF, (c) GAN, and (e) VAE: 50 wind power scenarios (grey) of a randomly selected day of the testing set along with the ten % (blue), 50 % (black), and 90 % (green) quantiles, and the observations (red). Right part (b) NF, (d) GAN, and (f) VAE: the corresponding Pearson time correlation matrices of these scenarios with the periods as rows and columns. The NF tends to exhibit no time correlation between scenarios. In contrast, the VAE and GAN tend to be partially time-correlated over a few periods.

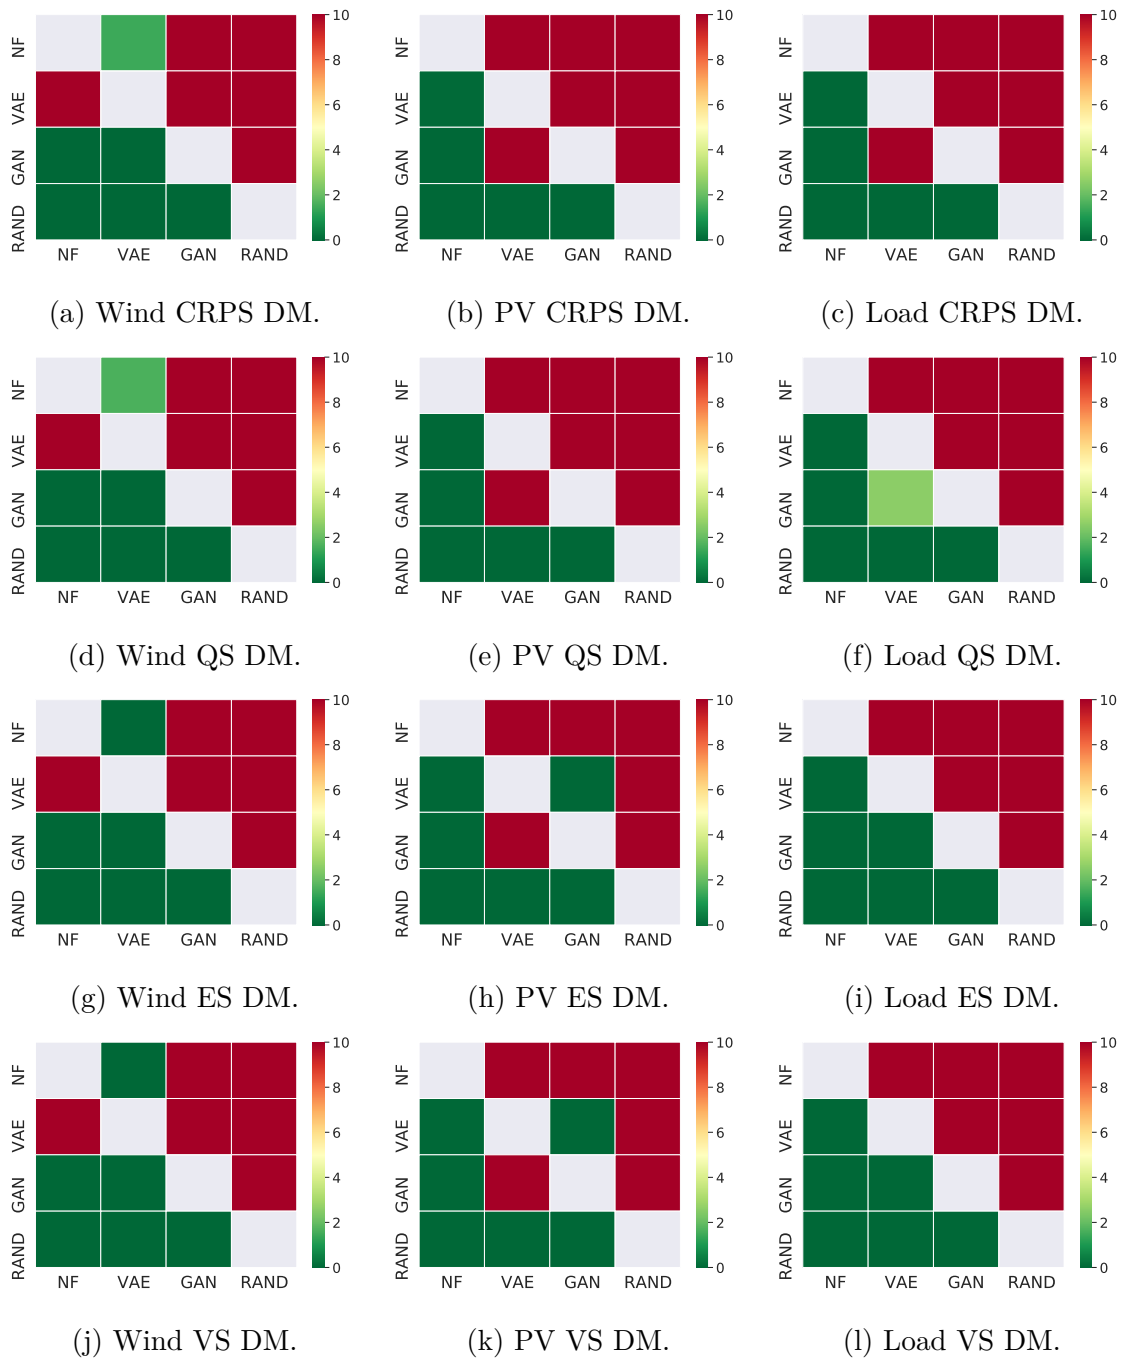


Fig. 6.13 Wind, PV and load tracks Diebold-Mariano tests.

Wind track: the DM tests of the CRPS, QS, ES, and VS confirm that the VAE outperforms the NF for these metrics. PV and load tracks: the NF outperforms the VAE and GAN for all metrics. Note: the GAN outperforms the VAE for both the ES and VS for the PV track. However, the VAE outperforms the GAN on this dataset for both the CRPS and QS.




Fig. 6.14 PV and load scenarios shape comparison and analysis.

For each model, there are 50 PV and load scenarios (grey) of a randomly selected day of the testing set along with the 10 % (blue), 50 % (black), and 90 % (green) quantiles, and the observations (red). The corresponding Pearson time correlation matrices of these scenarios with the periods as rows and columns are displayed. Like wind power scenarios, NF tends to exhibit no time correlation between PV and load tracks scenarios. In contrast, the VAE and GAN tend to be partially time-correlated over a few periods for the PV track. Furthermore, the VAE and GAN tend to be highly time-correlated for the load track.




Fig. 6.15 Average of the correlation matrices over the testing set for the three datasets. Left: wind power; center: PV; right: load. The trend in terms of time correlation is observed on each day of the testing set for all tracks. The NF scenarios are not correlated. In contrast, the VAE and GAN scenarios tend to be time-correlated over a few periods. In particular, the VAE generates highly time-correlated scenarios for the load dataset.

6.6 Conclusions and perspectives

This Chapter proposed a fair and thorough comparison of NFs with the state-of-the-art deep learning generative models, GANs and VAEs, both in quality and value. The numerical experiments employ the open data of the Global Energy Forecasting Competition 2014. The generative models use the conditional information to compute improved weather-based PV, wind power, and load scenarios. This Chapter demonstrated that NFs can challenge GANs and VAEs as they are, overall, more accurate both in terms of quality and value (see Chapter 13). In addition, they can be used effectively by non-expert deep learning practitioners. In addition, NFs have several advantages over more traditional deep learning approaches that should motivate their introduction into power system applications:

1. NFs directly learn the stochastic multivariate distribution of the underlying process by maximizing the likelihood. Therefore, in contrast to VAEs and GANs, NFs provide access to the exact likelihood of the model's parameters, hence offering a sound and direct way to optimize the network parameters [177]. It may open a new range of advanced applications benefiting from this advantage. For instance, to transfer scenarios from one location to another based on the knowledge of the probability density function. A second application is the importance sampling for stochastic optimization based on a scenario approach. Indeed, NFs provide the likelihood of each generated scenario, making it possible to filter relevant scenarios in stochastic optimization.
2. In our opinion, NFs are easier to use by non-expert deep learning practitioners once the libraries are available, as they are more reliable and robust in terms of hyper-parameters selection. GANs and VAEs are particularly sensitive to the latent space dimension, the structure of the neural networks, the learning rate, *etc.* GANs convergence, by design, is unstable, and for a given set of hyper-parameters, the scenario's quality may differ completely. In contrast, it was easier to retrieve relevant NFs hyper-parameters by manually testing a few sets of values, satisfying training convergence, and quality results.

Nevertheless, their usage as a base component of the machine learning toolbox is still limited compared to GANs or VAEs.

6.7 Appendix: Table 6.1 justifications

Wang et al. [175] use a Wasserstein GAN with gradient penalty to model both the uncertainties and the variations of the load. First, point forecasting is conducted, and the corresponding residuals are derived. Then, the GAN generates residual scenarios conditional on the day type, temperatures, and historical loads. The GAN model is compared with the same version without gradient penalty and two quantile regression models: random forest and gradient boosting regression tree. The quality evaluation is conducted on open load datasets from the Independent System Operator-New England³ with five metrics: (1) the continuous ranked probability score; (2) the quantile score; (3) the Winkler score; (4) reliability diagrams; (5) Q-Q plots. Note: the forecast value is not assessed.

Qi et al. [147] propose a concentrating solar power (CSP) configuration method to determine the CSP capacity in multi-energy power systems. The configuration model considers the uncertainty by scenario analysis. A β VAE generates the scenarios. It is an improved version of the original VAE. However, it does not consider weather forecasts, and the model is trained only by using historical observations. The quality evaluation is conducted on two wind farms and six PV plants using three metrics. (1) The leave-one-out accuracy of the 1-nearest neighbor classifier. (2) The comparison of the frequency distributions of the actual data and the generated scenarios. (3) Comparing the spatial and temporal correlations of the actual data and the scenarios by computing Pearson correlation coefficients. The value is assessed by considering the case study of the CSP configuration model, where the β VAE is used to generate PV, wind power, and load scenarios. However, the VAE is not compared to another generative model for both the quality and value evaluations. Note: the dataset does not seem to be in open-access. Finally, the value evaluation case study is not trivial due to the mathematical formulation that requires a certain level of knowledge of the system. Thus, the replicability criterion is partially satisfied.

Ge et al. [63] compared NFs to VAEs and GANs for the generation of daily load profiles. The models do not take into account weather forecasts but only historical observations. However, an example is given to illustrate the principle of generating conditional daily load profiles by using three groups: light load, medium load, and heavy load. The quality evaluation uses five indicators. Four to assess the temporal

³<https://www.iso-ne.com/>

correlation: (1) probability density function; (2) autocorrelation function; (3) load duration curve; (4) a wave rate is defined to evaluate the volatility of the daily load profile. Furthermore, one additional for the spatial correlation: (5) Pearson correlation coefficient is used to measure the spatial correlation among multiple daily load profiles. The simulations use the open-access London smart meter and Spanish transmission service operator datasets of Kaggle. The forecast value is not assessed.

6.8 Appendix: background

6.8.1 NFs

NF computation

Evaluating the likelihood of a distribution modeled by a normalizing flow requires computing (6.2), *i.e.*, the normalizing direction, as well as its log-determinant. Increasing the number of sub-flows by K of the transformation results in only $\mathcal{O}(K)$ growth in the computational complexity as the log-determinant of J_{f_θ} can be expressed as

$$\log |\det J_{f_\theta}(\mathbf{x})| = \log \left| \prod_{k=1}^K \det J_{f_{k,\theta}}(\mathbf{x}) \right|, \quad (6.16a)$$

$$= \sum_{k=1}^K \log |\det J_{f_{k,\theta}}(\mathbf{x})|. \quad (6.16b)$$

However, with no further assumption on f_θ , the computational complexity of the log-determinant is $\mathcal{O}(K \cdot T^3)$, which can be intractable for large T . Therefore, the efficiency of these operations is essential during training, where the likelihood is repeatedly computed. There are many possible implementations of NFs detailed by Papamakarios et al. [137], Kobyzev et al. [105] to address this issue.

Autoregressive flow

The Jacobian of the autoregressive transformation f_θ defined by (6.3) is lower triangular, and its log-absolute-determinant is

$$\log |\det J_{f_\theta}(\mathbf{x})| = \log \prod_{i=1}^T \left| \frac{\partial f^i}{\partial x_i}(x_i; h^i) \right|, \quad (6.17a)$$

$$= \sum_{i=1}^T \log \left| \frac{\partial f^i}{\partial x_i}(x_i; h^i) \right|, \quad (6.17b)$$




Fig. 6.16 VAE process. Credits: Francois Fleuret, EE559 Deep Learning, EPFL <https://fleuret.org/dlc/>

that is calculated in $\mathcal{O}(T)$ instead of $\mathcal{O}(T^3)$.

Affine autoregressive flow

A simple choice of transformer is the class of affine functions

$$f^i(x_i; h^i) = \alpha_i x_i + \beta_i, \quad (6.18)$$

where $f^i(\cdot; h^i) : \mathbb{R} \rightarrow \mathbb{R}$ is parameterized by $h^i = \{\alpha_i, \beta_i\}$, α_i controls the scale, and β_i controls the location of the transformation. Invertibility is guaranteed if $\alpha_i \neq 0$, and this can be easily achieved by *e.g.* taking $\alpha_i = \exp(\tilde{\alpha}_i)$, where $\tilde{\alpha}_i$ is an unconstrained parameter in which case $h^i = \{\tilde{\alpha}_i, \beta_i\}$. The derivative of the transformer with respect to x_i is equal to α_i . Hence the log-absolute-determinant of the Jacobian becomes

$$\log |\det J_{f_\theta}(\mathbf{x})| = \sum_{i=1}^T \log |\alpha_i| = \sum_{i=1}^T \tilde{\alpha}_i. \quad (6.19)$$

Affine autoregressive flows are simple and computation efficient. However, they are limited in expressiveness requiring many stacked flows to represent complex distributions. It is unknown whether affine autoregressive flows with multiple layers are universal approximators [137].

6.8.2 VAEs

Figure 6.16 illustrates the VAE process with f the encoder and g the decoder.

Gradients computation

By using (6.6) $\mathcal{L}_{\theta,\varphi}$ is decomposed in two parts

$$\mathcal{L}_{\theta,\varphi}(\mathbf{x}, \mathbf{c}) = \mathbb{E}_{q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})] - \text{KL}[q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p(\mathbf{z})]. \quad (6.20)$$

$\nabla_\theta \mathcal{L}_{\theta,\varphi}$ is estimated with the usual Monte Carlo gradient estimator. However, the estimation of $\nabla_\varphi \mathcal{L}_{\theta,\varphi}$ requires the reparameterization trick proposed by Kingma and Welling [102], where the random variable \mathbf{z} is re-expressed as a deterministic variable

$$\mathbf{z} = g_\varphi(\epsilon, \mathbf{x}), \quad (6.21)$$

with ϵ an auxiliary variable with independent marginal p_ϵ , and $g_\varphi(\cdot)$ some vector-valued function parameterized by φ . Then, the first right hand side of (6.20) becomes

$$\mathbb{E}_{q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})] = \mathbb{E}_{p(\epsilon)} [\log p_\theta(\mathbf{x}|g_\varphi(\epsilon, \mathbf{x}), \mathbf{c})]. \quad (6.22)$$

$\nabla_\varphi \mathcal{L}_{\theta,\varphi}$ is now estimated with Monte Carlo integration.

Conditional VAE implemented

Following Kingma and Welling [102], we implemented Gaussian multi-layer perceptrons (MLPs) for both the encoder NN_φ and decoder NN_θ . In this case, $p(\mathbf{z})$ is a centered isotropic multivariate Gaussian, $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})$ and $q_\varphi(\mathbf{x}|\mathbf{z}, \mathbf{c})$ are both multivariate Gaussian with a diagonal covariance and parameters $\boldsymbol{\mu}_\theta, \boldsymbol{\sigma}_\theta$ and $\boldsymbol{\mu}_\varphi, \boldsymbol{\sigma}_\varphi$, respectively. Note that there is no restriction on the encoder and decoder architectures, and they could as well be arbitrarily complex convolutional networks. Under these assumptions, the conditional VAE implemented is

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I}), \quad (6.23a)$$

$$p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_\theta, \boldsymbol{\sigma}_\theta^2 \mathbf{I}), \quad (6.23b)$$

$$q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_\varphi, \boldsymbol{\sigma}_\varphi^2 \mathbf{I}), \quad (6.23c)$$

$$\boldsymbol{\mu}_\theta, \log \boldsymbol{\sigma}_\theta^2 = \text{NN}_\theta(\mathbf{x}, \mathbf{c}), \quad (6.23d)$$

$$\boldsymbol{\mu}_\varphi, \log \boldsymbol{\sigma}_\varphi^2 = \text{NN}_\varphi(\mathbf{z}, \mathbf{c}). \quad (6.23e)$$

Then, by using the valid reparameterization trick proposed by Kingma and Welling [102]

$$\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6.24a)$$

$$\mathbf{z} := \boldsymbol{\mu}_\varphi + \boldsymbol{\sigma}_\varphi \boldsymbol{\epsilon}, \quad (6.24b)$$

$\mathcal{L}_{\theta, \varphi}$ is computed and differentiated without estimation using the expressions

$$\text{KL}[q_\varphi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p(\mathbf{z})] = -\frac{1}{2} \sum_{j=1}^d (1 + \log \boldsymbol{\sigma}_{\varphi, j}^2 - \boldsymbol{\mu}_{\varphi, j}^2 - \boldsymbol{\sigma}_{\varphi, j}^2), \quad (6.25a)$$

$$\mathbb{E}_{p(\boldsymbol{\epsilon})} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})] \approx -\frac{1}{2} \left\| \frac{\mathbf{x} - \boldsymbol{\mu}_\theta}{\boldsymbol{\sigma}_\theta} \right\|^2, \quad (6.25b)$$

with d the dimensionality of \mathbf{z} .

6.8.3 GANs

GAN

The original GAN value function $V(\phi, \theta)$ proposed by Goodfellow et al. [74] is

$$V(\phi, \theta) = \underbrace{\mathbb{E}_{\mathbf{x}}[\log d_\phi(\mathbf{x}|\mathbf{c})] + \mathbb{E}_{\hat{\mathbf{x}}}[\log(1 - d_\phi(\hat{\mathbf{x}}|\mathbf{c}))]}_{:= -L_d}, \quad (6.26a)$$

$$L_g := -\mathbb{E}_{\hat{\mathbf{x}}}[\log(1 - d_\phi(\hat{\mathbf{x}}|\mathbf{c}))], \quad (6.26b)$$

where L_d is the cross-entropy, and L_g the probability the discriminator wrongly classifies the samples.

WGAN

The divergences which GANs typically minimize are responsible for their training instabilities for reasons investigated theoretically by Arjovsky and Bottou [7]. Arjovsky et al. [8] proposed instead using the *Earth mover* distance, also known as the Wasserstein-1 distance

$$W_1(p, q) = \inf_{\gamma \in \Pi(p, q)} \mathbb{E}_{(x, y) \sim \gamma} [\|x - y\|], \quad (6.27)$$

where $\Pi(p, q)$ denotes the set of all joint distributions $\gamma(x, y)$ whose marginals are respectively p and q , $\gamma(x, y)$ indicates how much mass must be transported from x

to y in order to transform the distribution p into q , $\|\cdot\|$ is the L1 norm, and $\|x - y\|$ represents the cost of moving a unit of mass from x to y . However, the infimum in (6.27) is intractable. Therefore, Arjovsky et al. [8] used the Kantorovich-Rubinstein duality [170] to propose the Wasserstein GAN (WGAN) by solving the min-max problem

$$\theta^* = \arg \min_{\theta} \max_{\phi \in \mathcal{W}} \mathbb{E}_{\mathbf{x}}[d_{\phi}(\mathbf{x}|\mathbf{c})] - \mathbb{E}_{\hat{\mathbf{x}}}[d_{\phi}(\hat{\mathbf{x}}|\mathbf{c})], \quad (6.28)$$

where $\mathcal{W} = \{\phi : \|d_{\phi}(\cdot)\|_L \leq 1\}$ is the 1-Lipschitz space, and the classifier $d_{\phi}(\cdot) : \mathbb{R}^T \times \mathbb{R}^{|\mathbf{c}|} \rightarrow [0, 1]$ is replaced by a critic function $d_{\phi}(\cdot) : \mathbb{R}^T \times \mathbb{R}^{|\mathbf{c}|} \rightarrow \mathbb{R}$. However, the weight clipping used to enforce d_{ϕ} 1-Lipschitzness can lead sometimes the WGAN to generate only poor samples or failure to converge [79]. Therefore, we implemented the WGAN-GP to tackle this issue.

Chapter 7

Part I conclusions

“If you’re lonely when you’re alone, you’re in bad company.”

— **Jean-Paul Sartre**

Part I presents the forecasting tools and metrics required to produce and evaluate reliable point and probabilistic forecasts used as input of decision-making models. The forecasts take various forms: point forecasts, quantiles, prediction intervals, density forecasts, and scenarios. An example of forecast quality evaluation is given with standard deep learning models such as recurrent neural networks that compute PV and electrical consumption point forecasts on a real case study. Then, several case studies are used to assess the forecast quality of probabilistic forecasts.

- Deep learning models such as the encoder-decoder architecture and recurrent neural networks compute PV quantile forecasts. These models are trained by quantile regression. The forecast quality is evaluated on a real case study composed of the PV generation of the parking rooftops of the Liège University. The quantile regression models are compared to quantiles derived from deep learning generative models. In terms of forecast quality, the latter models outperform, in this case, the quantile regression models. The forecast value is assessed in Part II where the quantile forecasts are used as input of robust optimization planners in the form of prediction intervals.
- A density forecast-based approach produces probabilistic forecasts of imbalance prices, focusing on the Belgian case. The two-step probabilistic approach computes the net regulation volume state transition probabilities. Then, it infers the imbalance price. A numerical comparison of this approach to standard forecasting

techniques is performed on the Belgium case. The proposed model outperforms other approaches on probabilistic error metrics but is less accurate at predicting the imbalance prices. Deep learning models could improve this probabilistic approach to avoid simplifying assumptions and adding input features to describe the market situation better.

- Finally, a recent class of deep learning generative models, the normalizing flows, is investigated. A fair and thorough comparison is conducted with the state-of-the-art deep learning generative models, generative adversarial networks and variational autoencoders. The experiments employ the open data of the Global Energy Forecasting Competition 2014. The models use the conditional information to compute improved weather-based PV, wind power, and load scenarios. The results demonstrate that normalizing flows can challenge generative adversarial networks and variational autoencoders. Indeed, they are overall more accurate both in terms of quality and value (see Chapter 13). Furthermore, they can be used effectively by non-expert deep learning practitioners.

Part II

Planning and control

Overview

Part II presents the decision-making tools under uncertainty required to address the day-ahead planning of a microgrid. It investigates several approaches: deterministic planning using linear programming based on point forecasts, stochastic programming using scenarios, and robust optimization using quantile forecasts. Several case studies are considered: a grid-connected microgrid using a value function approach to propagate the information from planning to real-time optimization, a grid-connected microgrid in the capacity firming framework, and an energy retailer on the day-ahead market. The case studies use the forecasting techniques studied in Part I.

“*The straight line, a respectable optical
illusion which ruins many a man.*”

— **Victor Hugo**

Figure 7.1 illustrates the organization of the second part. It addresses the energy management of a grid-connected microgrid by using the forecasting techniques studied in Part I. Chapter 8 provides the optimization basics. Chapter 9 proposes a value function-based approach as a way to propagate information from operational planning to real-time optimization. Chapters 10, 11, and 12 propose stochastic, sizing, and robust approaches to address the energy management of a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market, respectively. The capacity firming framework has been designed for isolated markets, such as the Overseas France islands. Chapter 13 is the extension of Chapter 6. It investigates the forecast value assessment of the deep generative models by considering an energy retailer portfolio on the day-ahead market. Finally, Chapter 14 draws the conclusions of Part II.




Fig. 7.1 Part II skeleton.

Chapter 8

Decision-making background

Overview

This Chapter introduces some basics of linear programming and optimization methodologies to address uncertainty in decision-making used to formulate the problems considered in Part II.

The interested reader is referred to more general textbooks for further information [124, 16, 22] and the lectures of the courses "Renewables in Electricity Markets"^a and "Advanced Optimization and Game Theory for Energy Systems"^b given by professor Pierre Pinson and associate professor Jalal Kazempour, respectively, at the Technical University of Denmark.

^a<http://pierrepinson.com/index.php/teaching/>

^b<https://www.jalalkazempour.com/teaching>

“ Nothing is more difficult, and therefore more precious, than to be able to decide.”

— **Napoleon Bonaparte**

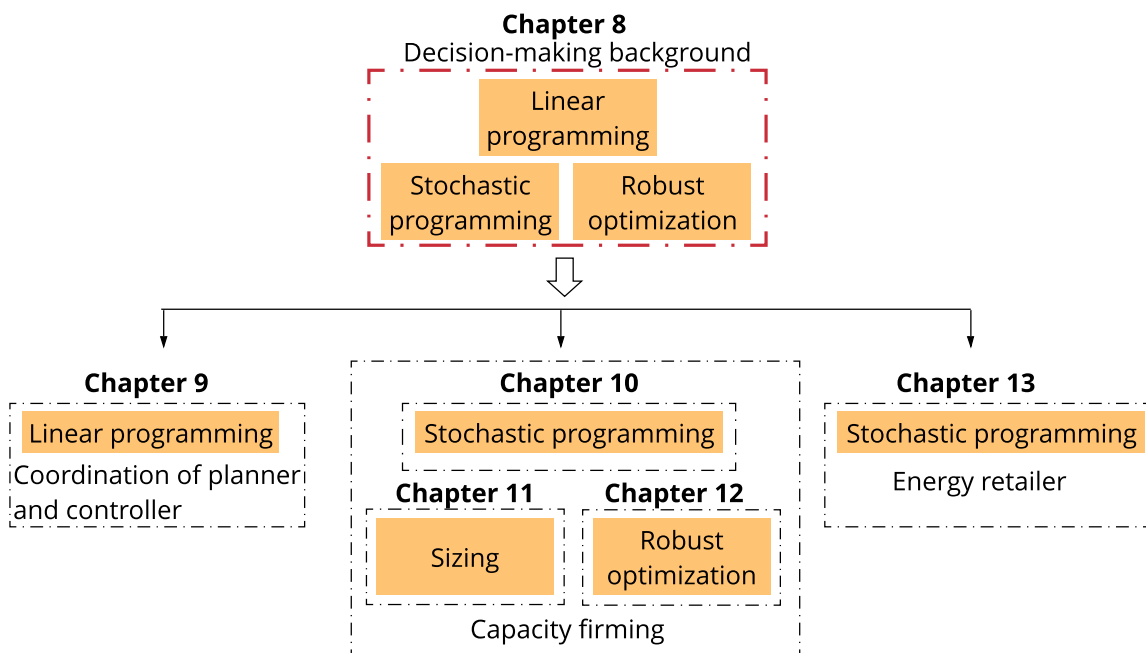


Fig. 8.1 Chapter 8 position in Part II.

Section 8.1 introduces the mathematical formulation of a linear optimization problem. Sections 8.2 and 8.3 present the mathematical formulations of a linear optimization problem when considering uncertainty in the parameters of the problem. First, by considering a stochastic programming approach. Second, by using a robust methodology.

8.1 Linear programming

8.1.1 Formulation of a linear programming problem

The most straightforward instance of an optimization problem is a linear programming problem.

Definition 8.1.1 (Linear programming). Linear programming (LP) minimizes a linear cost function subject to linear equality and inequality constraints.

In a LP problem, we are given a cost vector $\mathbf{c} \in \mathbb{R}^n$ and we seek to minimize a linear cost function $\mathbf{c}^\top \mathbf{x}$ over all n -dimensional vectors $\mathbf{x} \in \mathbb{R}^n$ subject to a set of linear equality and inequality constraints. Suppose that there is a total of m such constraints and let $\mathbf{b} \in \mathbb{R}^m$ and \mathbf{A} be the $m \times n$ matrix. The constraints can be expressed compactly in the form $\mathbf{Ax} \geq \mathbf{b}$.

Definition 8.1.2 (LP in standard form). [16, Chapter 1] A linear programming problem of the form

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{Ax} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{8.1}$$

is said to be in the *standard form*.

LP problems model a wide variety of real-world problems, such as day-ahead planning of a microgrid.

8.1.2 Duality in linear programming

Duality theory deals with the relation between a LP problem, called the *primal*, and another LP problem, called the *dual*. It uncovers the deeper structure of a LP problem and is a powerful tool that has various applications. The duality theory is motivated by the Lagrange multiplier method, where a price variable is associated with each

constraint. It explores prices under which the presence or absence of the constraints does not affect the optimal cost. Then, the correct prices can be found by solving a new LP problem, called the original's dual [16].

Definition 8.1.3 (Dual of a LP problem). [16, Chapter 4] Given a primal problem in the standard form (8.1), its dual is defined as

$$\begin{aligned} \max_{\mathbf{p}} \quad & \mathbf{p}^\top \mathbf{b} \\ \text{s.t.} \quad & \mathbf{p}^\top \mathbf{A} \leq \mathbf{c}^\top, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{8.2}$$

with \mathbf{p} the price vector of the same dimension as \mathbf{b} .

Note that the dual of the dual problem is the primal problem. The *strong duality* theorem is the central result of linear programming duality.

'If a linear programming problem has an optimal solution so does its dual, and the respective optimal costs are equal.' [16, Chapter 4, Theorem 4.4]

8.2 Stochastic optimization

Most decision-making problems are subject to uncertainty due to the intrinsic stochasticity of natural events conditioning our choices, *e.g.*, the weather, or, more generally, to the inaccurate knowledge of input information. Therefore, decision-makers are interested in methods and tools that provide solutions less sensitive to environmental influences or inaccurate data while simultaneously reducing cost, increasing profit, or improving reliability [124, Appendix C].

Let consider a LP in the standard form (8.1). If parameters \mathbf{A} , \mathbf{b} , and \mathbf{c} are perfectly known, solution algorithms for linear optimization problems, *e.g.*, the simplex method, can be used to find the best value of the decision variable vector \mathbf{x} .

However, suppose some of these parameters are contingent on the realization $\xi(\omega)$ of a particular random vector ξ . In that case, determining the optimal solution to the problem (8.1) may become further challenging. First, the issue of how to guarantee the feasibility of decision vector \mathbf{x} becomes remarkably more involved when optimizing under uncertainty because \mathbf{A} and \mathbf{b} are not completely known in advance. Second, the issue of how to guarantee the optimality of decision vector \mathbf{x} is at stake because \mathbf{c} is not completely known in advance. Finally, the LP problem (8.1) needs to be recast so that

solution algorithms for linear programming problems can be used to obtain the optimal value of decision vector \mathbf{x} taking into account the uncertainty of the parameters.

Stochastic programming provides the concepts and tools required to deal with the implications of having uncertain data in an optimization problem for decision making to address the three issues previously stated. It assumes that an accurate probabilistic description of the random variable is assumed available, under the form of the probability distributions or densities. In the following, we consider stochastic programming with *recourse* where the set of decisions is divided into two groups: (1) decisions have to be taken before the realization of uncertain parameters. These decisions are known as *first-stage*, \mathbf{x} , or here-and-now decisions and do not depend on the realization of the random parameters; (2) decisions can be taken after the realization of uncertain parameters. These decisions are called *second-stage*, $\mathbf{y}(\mathbf{x}, \omega)$, or recourse decisions and are dependent on each plausible value of the random parameters. Note: the term recourse points to the fact that second-stage decisions enable the decision-maker to adapt to the actual outcomes of the random events.

Definition 8.2.1 (Two-stage program with fixed recourse). [22, Chapter 2] A classical two-stage stochastic linear program with fixed recourse is

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} + \mathbb{E}_{\xi} [Q(\mathbf{x}, \xi(\omega))] \\ \text{s.t.} \quad & \mathbf{Ax} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{8.3}$$

with

$$\begin{aligned} Q(\mathbf{x}, \xi(\omega)) = \min_{\mathbf{y}} \quad & \mathbf{q}(\omega)^\top \mathbf{y} \\ \text{s.t.} \quad & \mathbf{Wy} = \mathbf{h}(\omega) - T(\omega)\mathbf{x}, \quad \forall \omega \in \Omega \\ & \mathbf{y} \geq \mathbf{0}, \quad \forall \omega \in \Omega \end{aligned} \tag{8.4}$$

the second-stage value function.

Note: the uncertainty involved in problem (8.3) and (8.4) is assumed to be properly represented by means of a finite set Ω of scenarios ω , with a probability α_ω such that $\sum_{\omega \in \Omega} \alpha_\omega = 1$.

Definition 8.2.2 (Deterministic equivalent problem). [22, Chapter 2] The *deterministic equivalent* problem of the stochastic programming problem (8.3)-(8.4) is

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{y}} \quad & \mathbf{c}^\top \mathbf{x} + \sum_{\omega \in \Omega} \alpha_\omega \mathbf{q}(\omega)^\top \mathbf{y} \\ \text{s.t.} \quad & \mathbf{Ax} = \mathbf{b}, \\ & \mathbf{Wy} = \mathbf{h}(\omega) - \mathbf{T}(\omega)\mathbf{x}, \quad \forall \omega \in \Omega \\ & \mathbf{x} \geq \mathbf{0}, \\ & \mathbf{y} \geq \mathbf{0}, \quad \forall \omega \in \Omega \end{aligned} \tag{8.5}$$

and can be directly processed by off-the-shelf optimization software for linear programs.

8.3 Robust optimization

Robust optimization considers optimization problems with uncertain parameters not modeled using probability distributions but with uncertainty sets. A robust optimization strategy investigates a solution to an optimization problem that is feasible for any realization of the uncertain parameters within the uncertainty set and optimal for the worst-case realization of these uncertain parameters [124, Appendix D].

Definition 8.3.1 (Two-stage robust optimization with fixed recourse). The general form of two-stage robust optimization formulation is

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} + \max_{\mathbf{u} \in \mathcal{U}} \min_{\mathbf{y} \in \Omega(\mathbf{x}, \mathbf{u})} \mathbf{q}^\top \mathbf{y} \\ \text{s.t.} \quad & \mathbf{Ax} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{8.6}$$

with \mathcal{U} the uncertainty set and

$$\Omega(\mathbf{x}, \mathbf{u}) = \{\mathbf{y} \geq \mathbf{0} : \mathbf{Wy} = \mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{Mu}\}. \tag{8.7}$$

The objective of the problem (8.6) is to make the best decisions represented by variable vector \mathbf{x} for the worst realization of parameters in vector \mathbf{u} and considering the recourse decisions described by variable vector \mathbf{y} . If the right-hand-side problem, the \mathbf{y} -problem, is convex, it can be replaced by its dual and merged with the middle \mathbf{u} -problem rendering it a conventional single-level maximization problem. Overall, the resulting problem is a min-max problem that in some cases can be solved using decomposition

such as the Benders-dual cutting plane method [18] or column-and-constraint generation algorithm [183]. Both the Benders-dual method and the column-and-constraint generation procedure are implemented in a master sub-problem framework.

8.3.1 Benders-dual cutting plane algorithm

The key idea of the Benders-dual cutting plane (BD) algorithm is to gradually construct the value function of the first-stage decisions using dual solutions of the second-stage decision problems.

Consider the case where the second-stage decision problem is a linear programming (LP) problem in \mathbf{x} . We first take the *relatively complete recourse* assumption that this LP is feasible for any given \mathbf{x} and \mathbf{u} . Let \mathbf{p} be its dual variables. Then, we obtain its dual problem, which is a maximization problem and can be merged with the maximization over \mathbf{u} . As a result, we have the following dispatch problem, which yields the sub-problem (SP) in the BD algorithm.

Definition 8.3.2 (BD sub-problem).

$$\begin{aligned} R(\mathbf{x}) = \max_{\mathbf{u}, \mathbf{p}} & (\mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{M}\mathbf{u})^\top \mathbf{p} \\ \text{s.t.} & \mathbf{W}^\top \mathbf{p} \geq \mathbf{q}, \\ & \mathbf{u} \in \mathcal{U}. \end{aligned} \quad (8.8)$$

Note: that the resulting problem in (8.8) is a bilinear optimization problem. Several approaches have been developed to address this issue and depend on the case study.

Suppose, we have an oracle that can solve (8.8) for a given first stage variable \mathbf{x} . Then, the optimal solution is $(\mathbf{p}^*, \mathbf{u}^*)$, and a cutting plane in the form of

$$\theta \geq (\mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{M}\mathbf{u}^*)^\top \mathbf{p}^* \quad (8.9)$$

can be generated, and included in the master problem (MP).

Definition 8.3.3 (BD master problem). At iteration k of the BD master problem is

$$\min_{\mathbf{x}, \theta} \mathbf{c}^\top \mathbf{x} + \theta \quad (8.10)$$

$$\text{s.t. } \mathbf{A}\mathbf{x} = \mathbf{b},$$

$$\theta \geq (\mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{M}\mathbf{u}_l^*)^\top \mathbf{p}_l^* \quad \forall l \leq k \quad (8.11)$$

$$\mathbf{x} \geq \mathbf{0}, \quad \theta \in \mathbb{R},$$

which can compute an optimal first stage solution \mathbf{x}_k^* .

The MP and SP provide the lower and upper bounds with $\mathbf{c}^\top \mathbf{x}_k^* + \theta_k^*$ and $R(\mathbf{x}_{k-1}^*)$, respectively, at iteration k . When considering the relatively complete recourse assumption, [183, 18, 14] demonstrated the BD algorithm converges to the optimal solution of the two-stage robust optimization problem. The lower and upper bounds converge in a finite number of steps by iteratively introducing cutting planes (8.9) and computing the MP (8.10). Note: if the SP is a MILP, this result is not straightforward. It is the case in the capacity firming problem studied in Chapter 12 where this issue is addressed.

8.3.2 Column and constraints generation algorithm

The column and constraints generation (CCG) algorithm does not create constraints using dual solutions of the second-stage decision problem. It dynamically generates constraints with recourse decision variables in the primal space for an identified scenario. Let assume the uncertainty set \mathcal{U} is a finite discrete set with $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$ and $\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ are the corresponding recourse decision variables. Then, the 2-stage RO (8.6) can be reformulated as follows

$$\min_{\mathbf{x}, \theta} \mathbf{c}^\top \mathbf{x} + \theta \quad (8.12)$$

$$\text{s.t. } \mathbf{A}\mathbf{x} = \mathbf{b}, \quad (8.13)$$

$$\theta \geq \mathbf{q}^\top \mathbf{y}^l \quad \forall l \leq N \quad (8.14)$$

$$\mathbf{W}\mathbf{y}^l = \mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{M}\mathbf{u}^l \quad \forall l \leq N \quad (8.15)$$

$$\mathbf{x} \geq \mathbf{0}, \quad \mathbf{y}^l \geq \mathbf{0} \quad \forall l \leq N, \quad \theta \in \mathbb{R}. \quad (8.16)$$

Thus, it reduces to solve an equivalent, probably large-scale, MILP, which is very close to a 2-stage stochastic programming model if the probability distribution over \mathcal{U} is known. When the uncertainty set is large or is a polyhedron, enumerating all the possible uncertain scenarios in \mathcal{U} is not feasible. Constraints (8.14) indicate that not all scenarios, and their corresponding variables and constraints, are necessary for defining the optimal value \mathbf{x}^* of the 2-stage RO. It is likely that a few relevant scenarios, a small subset of the uncertainty set, play a significant role in the formulation. Therefore, the key idea of the CCG procedure is to generate recourse decision variables for the significant scenarios.

Definition 8.3.4 (CCG sub-problem). Similar to the BD method, the CCG algorithm uses a master sub-problem framework. Let assume an oracle can solve the dispatch

problem for a given first-stage \mathbf{x}

$$Q(\mathbf{x}) = \max_{\mathbf{u} \in \mathcal{U}} \min_{\mathbf{y} \in \Omega(\mathbf{x}, \mathbf{u})} \mathbf{q}^\top \mathbf{y}. \quad (8.17)$$

Then, the optimal solution \mathbf{u}^* can be derived to build constraints and variables in the master problem.

Definition 8.3.5 (CCG master problem). At iteration k of the CCG master problem is

$$\min_{\mathbf{x}, \theta} \mathbf{c}^\top \mathbf{x} + \theta \quad (8.18)$$

$$\text{s.t. } \mathbf{Ax} = \mathbf{b}, \quad (8.19)$$

$$\theta \geq \mathbf{q}^\top \mathbf{y}^l \quad \forall l \leq k \quad (8.20)$$

$$\mathbf{Wy}^l = \mathbf{h} - \mathbf{T}\mathbf{x} - \mathbf{M}\mathbf{u}^{*,l} \quad \forall l \leq k \quad (8.21)$$

$$\mathbf{x} \geq \mathbf{0}, \quad \mathbf{y}^l \geq \mathbf{0} \quad \forall l \leq k, \quad \theta \in \mathbb{R}, \quad (8.22)$$

which can compute an optimal first stage solution \mathbf{x}_k^* .

Similar to BD, the MP and SP provide the lower and upper bounds with $\mathbf{c}^\top \mathbf{x}_k^* + \theta_k^*$ and $Q(\mathbf{x}_{k-1}^*)$, respectively, at iteration k .

8.4 Conclusions

This Chapter introduces the basics of linear programming and approaches to handle uncertainty in the parameters with the stochastic programming and robust approach. Depending on the application, each approach has its pros and cons. When considering stochastic programming, the number of scenarios needed to describe the most plausible outcomes of the uncertain parameters may be huge, leading to large-scale optimization problems that may become difficult to solve or intractable. In this case, a robust approach provides an alternative and compact manner to describe uncertain parameters. However, the robust counterpart optimization problem may be challenging to solve and requires a decomposition technique such as a Benders-dual cutting plane or column and constraints generation algorithm that is not trivial to implement numerically.

Chapter 9

Coordination of the planner and controller

Overview

This Chapter presents a two-layer approach with a value function to propagate information from operational planning to real-time optimization. The value function-based approach shares some similarities with the coordination scheme proposed in Kumar et al. [108], which is based on stochastic dual dynamic programming. This study brings new contributions:

1. The approach is tested by accounting for forecasting errors and high-resolution data monitored on-site corresponding to a "real-life" case.
2. The value function approach allows to deal with indeterminacy issues. When there are several optimal solutions to the upper-level problem, this is accounted for in the lower level part, and a bias term can be added to favor one type of behavior over another, *e.g.*, charge early.
3. This methodology is fully compatible with the energy markets as it can deal with imbalance, reserve, and dynamic selling/purchasing prices.

This study reports results on an industrial microgrid capable of on/off-grid operation. Generation and consumption forecasts are based on weather forecasts obtained with the MAR model [59].

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Selmane Dakir, Clément Liu, and Bertrand Cornélusse. Coordination of operational planning and real-time optimization in microgrids. *Electric Power Systems Research*, 190:106634, 2021. URL <https://arxiv.org/abs/2106.02374>.

“When making a decision of minor importance, I have always found it advantageous to consider all the pros and cons. In vital matters, however, such as the choice of a mate or a profession, the decision should come from the unconscious, from somewhere within ourselves. In the important decisions of personal life, we should be governed, I think, by the deep inner needs of our nature.”

— Sigmund Freud




Fig. 9.1 Chapter 9 position in Part II.

The hierarchical microgrid control levels divide a global microgrid control problem in time and space [134]. Control levels range from distributed device level controllers that run at a high frequency to centralized controllers optimizing market integration that run much less frequently. For computation time reasons, centralized controllers are often subdivided into two levels. Operational planning controllers that optimize decisions over a time horizon of one or several days but with a market period resolution, *e.g.*, 15 minutes. Real-time optimization controllers that deal with actions within the current market period. The coordination of these two levels is paramount to achieving the safest and most profitable operational management of microgrids. Microgrid control and management can be achieved in several ways. Control techniques and the principles of energy-storage systems are summarized in Palizban et al. [134]. A classification of microgrid control strategies into primary, secondary, and tertiary levels is done in Olivares et al. [131]. The two-level approach has been intensively studied. A double-layer coordinated control approach, consisting of the scheduling layer and the dispatch layer, is adopted in Jiang et al. [94]. The schedule layer provides an economical operation scheme, including state and power of controllable units based on the look-ahead multi-step optimization. In contrast, the dispatch layer follows the scheduling layer by considering power flow and voltage limits. A two-stage dispatch strategy for grid-connected systems is discussed in Wu et al. [179], where the first stage deals with the day-ahead schedule, optimizing capital and operational cost. At the same time, the lower level handles the rescheduling of the units for few hours ahead with a time resolution of 15 min. A two-stage control strategy for a PV BESS-ICE (Internal Combustion Engine) microgrid is implemented in Sachs and Sawodny [153]. Discrete Dynamic Programming is used in the first layer, while the second layer problem is posed as a Boundary Value Problem. An approach with a high-level deterministic optimizer running at a slow timescale, 15 min, coupled to a low-level stochastic controller running at a higher frequency, 1 min, is studied in Cominesi et al. [36]. A two-layer predictive energy management system for microgrids with hybrid energy storage systems consisting of batteries and supercapacitors is considered in Ju et al. [96]. This approach incorporates the degradation costs of the hybrid energy storage systems. A practical Energy Management System for isolated microgrid which considers the operational constraints of Distributed Energy Resources, active-reactive power balance, unbalanced system configuration, and loading, and voltage-dependent loads is studied in Solanki et al. [160]. A two-layer mixed-integer linear programming predictive control strategy was implemented and tested in simulation and experimentally in Polimeni et al. [146]. Finally, Moretti et al. [125] implemented a two-layer predictive management

strategy for an off-grid hybrid microgrid featuring controllable and non-controllable generation units and a storage system.

It is organized as follows. Section 9.1 summarizes the notation. Section 9.2 formulates the problem in an abstract manner. Section 9.3 introduces the novel two-level value function-based approach and the assumptions made. Section 9.4 describes the numerical tests. Section 9.5 reports the results. Conclusions are drawn in Section 9.6. The methodology used for forecasting is reported in Section 3.3 of Chapter 3.

9.1 Notation

Sets and indices

Name	Description
d	Device index.
t, t'	RTO and OP time periods indexes.
$\tau(t)$	Beginning of the next market period at time t .
$\mathcal{T}_i(t)$	Set of RTO time periods = $\{t, t + \Delta t, \dots, t + T_i\}$.
$\mathcal{T}_a^m(t)$	Set of OP time periods = $\{\tau(t), \tau(t) + \Delta\tau, \dots, \tau(t + T_a)\}$.
T_a, T_l	Time durations, with $T_a \ll T_l$
\mathcal{D}^k	Set of non-flexible loads ($k = \text{nfl}$), sheddable loads ($k = \text{she}$), steerable generators ($k = \text{ste}$), non-steerable generators ($k = \text{nst}$), storage devices ($k = \text{sto}$).

Parameters

Name	Description	Unit
Δt	Time delta between t and the market period.	minutes
$\Delta\tau$	Market period.	minutes
$\eta^{\text{cha}}, \eta^{\text{dis}}$	Charge and discharge efficiencies.	%
\bar{P}, \underline{P}	Maximum charging and discharging powers.	kW
$C_{d,t}^{\text{nfl}}$	Non-flexible power consumption.	kW
$C_{d,t}^{\text{she}}$	Flexible power consumption.	kW
$S_{d,t}^{\text{init}}$	Initial state of charge of battery d .	kWh
p_h	Maximum peak over the last twelve months.	kW
π^{P}	Yearly peak power cost.	€/kW
π_{OP}^{S}	Unitary revenue for providing reserve.	€/kW

π_{RTO}^s	Unitary RTO symmetric reserve penalty.	€/kW
$\pi_{d,t}^k$	Cost of load shedding ($k = she$), generating energy ($k = ste$), curtailing generation ($k = nst$).	€/kWh
$\gamma_{d,t}^{sto}$	Fee to use the battery d .	€/kWh
π_t^e, π_t^i	Energy prices of export and import.	€/kWh
π_d^i, π_n^i	Energy prices of day and night imports.	€/kWh
I^{cap}, E^{cap}	Maximum import and export limits.	kW
PV_p, C_p	PV and consumption capacities.	kW
S_p	Storage capacity.	kWh
\bar{S}, \underline{S}	Maximum and minimum battery capacities.	kWh

Forecasted or computed variables

Name	Description	Unit
a_t	Action at t .	-
a_t^m	Purely market related actions.	-
a_t^d	Set-points to the devices of the microgrid	-
$a_{d,t}^k$	Fraction of load shed ($k = she$), generation activated ($k = ste$), generation curtailed ($k = nst$) ($[0, 1]$).	-
$a_{d,t}^{cha}, a_{d,t}^{dis}$	Fraction of the maximum charging and discharging powers used for battery d ($[0, 1]$).	-
e_t^{gri}, i_t^{gri}	Energy export and import.	kWh
$\delta p_{t'}$	OP peak difference between peak at t' and p_h .	kW
$\delta p_{\tau(t-\Delta\tau), \tau(t)}$	RTO peak difference between peak at $\tau(t)$ and p_h .	kW
s_t^{TSO}	TSO symmetric reserve signal (0; 1).	-
r^{sym}	Symmetric reserve.	kW
Δr^{sym}	Reserve difference between OP and RTO.	kW
$r_{d,t'}^{s+}, r_{d,t'}^{s-}$	Upward and downward reserves of power available and provided by storage device d .	kW
$s_{d,t}$	State of charge of battery d .	kWh
s_t	Microgrid state at time t .	-
s_t^m	Information related to the current market position.	-
s_t^d	State of the devices.	-
v_t	The cost-to-go function.	k€
$\hat{\omega}$	Forecast of a random vector ω .	-

\bar{X}	Average of a variable X .	kW
X_{max}, X_{min}	Maximum and minimum of X .	kW
X_{std}	Standard deviation of X .	kW
c_E, c_p, c_t	Energy, peak and total costs.	k€
Δ_p	Peak increment.	kW
I_{tot}, E_{tot}	Total import and export.	MWh

9.2 Problem statement

A global microgrid control problem can be defined, for a given microgrid design and configuration, as operating a microgrid safely and in an economically efficient manner, by harvesting as much renewable energy as possible, operating the grid efficiently, optimizing the service to the demand side, and optimizing other side goals. We refine this definition below and start by making a few assumptions.

9.2.1 Assumptions

In this study, the control optimizes economic criteria, which are only related to active power. All devices are connected to the same electrical bus, which can be connected or disconnected from a public grid permanently or dynamically. Device-level controllers offer an interface to communicate their operating point and constraints, *e.g.*, maximum charge power as a function of the current state, and implement control decisions to reach power set-points. Fast load-frequency control, islanding management, as well as reactive power control are not in scope. The microgrid is a price taker in energy and reserve markets.

9.2.2 Formulation

Abstractly, a microgrid optimization problem can be formulated as follows

$$\min_{\mathbf{a}} \sum_{t \in \mathcal{T}_l} c(a_t, s_t, \omega_t) \quad (9.1a)$$

$$\text{s.t. } \forall t \in \mathcal{T}_l, s_{t+\Delta t} = f(a_t, s_t, \omega_t, \Delta t), \quad (9.1b)$$

$$s_t \in \mathcal{S}_t. \quad (9.1c)$$

A controller has to return a set of actions $a_t = (a_t^m, a_t^d)$ at any time t over the life of the microgrid (\mathcal{T}_l). Actions should be taken as frequently as possible to cope with the economic impact of the variability of the demand and generation sides, but not too

often to let transients vanish, *e.g.*, every few seconds. The time delta between action a_t and the next action taken is denoted by Δt , and is not necessarily constant. Some of these actions are purely market-related a_t^m , while other actions are communicated as set-points to the devices of the microgrid a_t^d . The state $s_t = (s_t^m, s_t^d)$ of the microgrid at time t is thus also made of two parts. (1) s_t^d represents the state of the devices, such as a storage system or a flexible load. (2) s_t^m gathers information related to the current market position, such as the nominated net position of the microgrid over the next market periods. The cost function c gathers all the economic criteria considered. The transition function f describes the physical and net position evolution of the system. At time instants $t \in \{\Delta\tau, 2\Delta\tau, \dots\}$, with $\Delta\tau$ the market period, some costs are incurred based on the value of some state variables, which are then reset for the following market period. This problem is challenging to solve since the system's evolution is uncertain, actions have long-term consequences, and are both discrete and continuous. Furthermore, although functions f and c are assumed time-invariant, they are generally non-convex and parameterized with stochastic variables ω_t .

9.3 Proposed method

In practice, solving the microgrid optimization problem above amounts, at every time t , to forecasting the stochastic variables $\omega_{\mathcal{T}_l(t)}$, then solving the problem¹

$$\mathbf{a}_{\mathcal{T}_l(t)}^* = \arg \min \sum_{t' \in \mathcal{T}_l(t)} c(a_{t'}, s_{t'}, \hat{\omega}_{t'}) \quad (9.2a)$$

$$\text{s.t. } \forall t' \in \mathcal{T}_l(t), s_{t'+\Delta t'} = f(a_{t'}, s_{t'}, \hat{\omega}_{t'}, \Delta t'), \quad (9.2b)$$

$$s_{t'} \in S_{t'}, \quad (9.2c)$$

and applying a_t^* (potentially changing $a_t^{m,*}$ at some specific moments only). Forecasts are valid only for a relatively near future, and optimizing over a long time would be incompatible with the real-time operation. Thus, this problem is approximated by cropping the lookahead horizon to $\mathcal{T}_a(t) \subset \mathcal{T}_l(t)$. However, market decisions must be refreshed much less frequently than set-points. We thus propose to further decompose the problem in an operational planning problem (OPP) for $\mathcal{T}_a^m(t)$ that computes

¹Which is here expressed as a deterministic problem for simplicity but should be treated as a stochastic problem in practice.

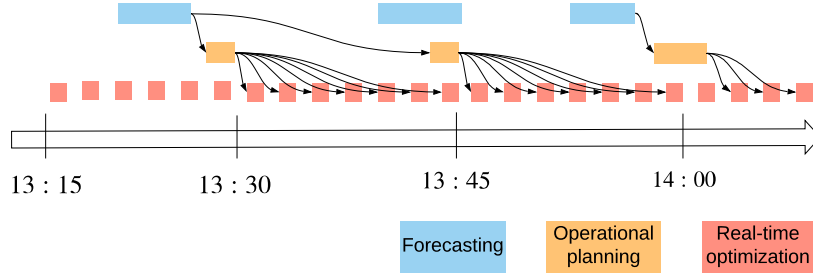


Fig. 9.2 Hierarchical control procedure illustration.

market decisions

$$\mathbf{a}_{\mathcal{T}_a^m(t)}^{m,*} = \arg \min \sum_{t' \in \mathcal{T}_a^m(t)} c^m(a_{t'}^m, s_{t'}, \hat{\omega}_{t'}) \quad (9.3a)$$

$$\text{s.t. } \forall t' \in \mathcal{T}_a^m(t), s_{t'+\Delta\tau} = f^m(a_{t'}^m, s_{t'}, \hat{\omega}_{t'}, \Delta\tau) \quad (9.3b)$$

$$s_{t'} \in S_{t'}, \quad (9.3c)$$

and a real-time problem (RTP) that computes set-points for time t

$$a_t^{d,*} = \arg \min c^d(a_t^d, s_t, \hat{\omega}_t) + v_{\tau(t)}(s_{\tau(t)}) \quad (9.4a)$$

$$\text{s.t. } s_{\tau(t)} = f^d(a_t^d, s_t, \hat{\omega}_t, \tau(t) - t) \quad (9.4b)$$

$$s_{\tau(t)} \in S_{\tau(t)}, \quad (9.4c)$$

with $c(a_t, s_t, w_t) = c^m(a_t^m, s_t, w_t) + c^d(a_t^d, s_t, w_t)$. The function v_t is the cost-to-go as a function of the system's state at the end of the ongoing market period. It regularizes decisions of RTP to account for the longer-term evolution of the system. We detail hereunder how we obtain v_t . An overview of the approach is depicted in Figure 9.2.

9.3.1 Computing the cost-to-go function

The function v_t represents the optimal value of (9.3) as a function of the initial state $s_{\tau(t)}$ of this problem. If we make the assumption that (9.3) is modeled as a linear program, the function $v_{\tau(t)}$ is thus convex and piecewise linear. Every evaluation of (9.3) with the additional constraint²

$$s_{\tau(t)} = s' \perp \mu \quad (9.5)$$

²The $\perp \mu$ notation means that μ is the dual variable of the constraint.

yields the value $v_{\tau(t)}(s')$ and a supporting inequality (a cut)

$$v_{\tau(t)}(s) \geq v_{\tau(t)}(s') + \mu^T s. \quad (9.6)$$

The algorithm to approximate $v_{\tau(t)}(s')$ works as follows:

1. estimate the domain of $v_{\tau(t)}$, *i.e.*, the range of states reachable at time $\tau(t)$ and the most probable state that will be reached $s_{\tau(t)}^*$;
2. evaluate $v_{\tau(t)}(s_{\tau(t)}^*)$ and the associated μ^* ;
3. repeat step 2 for other state values until all regions of $v_{\tau(t)}$ are explored.

Note: that if the state is of dimension one and (9.3) is a linear program, simplex basis validity information can be used to determine for which part of the domain of $v_{\tau(t)}$ the current cut is tight, else a methodology such as proposed in Bemporad et al. [10] can be used.

9.3.2 OPP formulation

The OPP objective function implemented for the case study is

$$J_{\mathcal{T}_a^m(t)}^{OP} = \sum_{t' \in \mathcal{T}_a^m(t)} \left(C_{t'}^{OP} + D_{t'}^{OP} \right) \quad (9.7)$$

with Operational Planner (OP) the name of this planer. $\mathcal{T}_a^m(t)$ is composed of 96 values with $\Delta\tau = 15$ minutes and $T_a = 24$ hours. $C_{t'}^{OP}$ models the immediate costs and $D_{t'}^{OP}$ the delayed costs at t' . $C_{t'}^{OP}$ takes into account different revenues and costs related to energy flows: the costs of shed demand, steered and non steered generation, the revenues from selling energy to the grid, the costs of purchasing energy from the grid and the costs for using storage

$$\begin{aligned} C_{t'}^{OP} = & \sum_{t' \in \mathcal{T}_a^m(t)} \left(\sum_{d \in \mathcal{D}^{\text{she}}} \Delta\tau \pi_{d,t'}^{\text{she}} C_{d,t'}^{\text{she}} a_{d,t'}^{\text{she}} + \sum_{d \in \mathcal{D}^{\text{ste}}} \Delta\tau \pi_{d,t'}^{\text{ste}} P_{d,t'}^{\text{ste}} a_{d,t'}^{\text{ste}} + \sum_{d \in \mathcal{D}^{\text{nst}}} \Delta\tau \pi_{d,t'}^{\text{nst}} P_{d,t'}^{\text{nst}} a_{d,t'}^{\text{nst}} \right. \\ & \left. + \sum_{d \in \mathcal{D}^{\text{sto}}} \Delta\tau \gamma_d^{\text{sto}} [\bar{P}_d \eta_d^{\text{cha}} a_{d,t'}^{\text{cha}} + \frac{P_d}{\eta_d^{\text{dis}}} a_{d,t'}^{\text{dis}}] - \pi_{t'}^e e_{t'}^{\text{gri}} + \pi_{t'}^i i_{t'}^{\text{gri}} \right). \end{aligned} \quad (9.8)$$

$D_{t'}^{OP}$ is composed of the peak cost and symmetric reserve revenue

$$D_{t'}^{OP} = \pi^p \delta p_{t'} - \pi_{OP}^s r_{t'}^{\text{sym}}, \quad (9.9)$$

$\delta p_{t'}$ is the peak difference between the previous maximum historic peak p_h and the current peak within the market period t' . $r_{t'}^{\text{sym}}$ is the symmetric reserve provided to the grid within the current market period t' .

9.3.3 OP constraints

The first set of constraints defines bounds on state and action variables, $\forall t' \in \mathcal{T}_a^m(t)$

$$a_{d,t'}^k \leq 1 \quad \forall d \in \mathcal{D}^k, \forall k \in \{\text{ste, she, nst}\} \quad (9.10a)$$

$$a_{d,t'}^{\text{cha}} \leq 1 \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.10b)$$

$$a_{d,t'}^{\text{dis}} \leq 1 \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.10c)$$

$$\underline{S}_d \leq s_{d,t'} \leq \bar{S}_d \quad \forall d \in \mathcal{D}^{\text{sto}}. \quad (9.10d)$$

The energy flows are constrained, $\forall t' \in \mathcal{T}_a^m(t)$, by

$$\begin{aligned} & (e_{t'}^{\text{gri}} - i_{t'}^{\text{gri}})/\Delta\tau - \sum_{d \in \mathcal{D}^{\text{nst}}} (1 - a_{d,t'}^{\text{nst}})P_{d,t'}^{\text{nst}} + \sum_{d \in \mathcal{D}^{\text{ste}}} a_{d,t'}^{\text{ste}}P_{d,t'}^{\text{ste}} \\ & + \sum_{d \in \mathcal{D}^{\text{nfl}}} C_{d,t'}^{\text{nfl}} + \sum_{d \in \mathcal{D}^{\text{she}}} (1 - a_{d,t'}^{\text{she}})C_{d,t'}^{\text{she}} + \sum_{d \in \mathcal{D}^{\text{sto}}} (\bar{P}_d a_{d,t'}^{\text{cha}} - \underline{P}_d a_{d,t'}^{\text{dis}}) = 0 \end{aligned} \quad (9.11a)$$

$$(e_{t'}^{\text{gri}} - i_{t'}^{\text{gri}})/\Delta\tau \leq E_{t'}^{\text{cap}} \quad (9.11b)$$

$$(i_{t'}^{\text{gri}} - e_{t'}^{\text{gri}})/\Delta\tau \leq I_{t'}^{\text{cap}}. \quad (9.11c)$$

The dynamics of the state of charge are, $\forall d \in \mathcal{D}^{\text{sto}}$

$$s_{d,1} - \Delta\tau(\bar{P}_d \eta_d^{\text{cha}} a_{d,1}^{\text{cha}} - \frac{P_d}{\eta_d^{\text{dis}}} a_{d,1}^{\text{dis}}) = S_d^{\text{init}} \quad (9.12a)$$

$$s_{d,t'} - s_{d,t'-\Delta\tau} - \Delta\tau(\bar{P}_d \eta_d^{\text{cha}} a_{d,t'}^{\text{cha}} - \frac{P_d}{\eta_d^{\text{dis}}} a_{d,t'}^{\text{dis}}) = 0, \quad \forall t' \in \mathcal{T}_a^m(t) \quad (9.12b)$$

$$s_{d,\tau(t+T_a)} = S_d^{\text{end}}. \quad (9.12c)$$

The set of constraints related to the peak power $\forall t' \in \mathcal{T}_a^m(t)$

$$(i_{t'}^{\text{gri}} - e_{t'}^{\text{gri}})/\Delta\tau \leq p_{t'} \quad (9.13a)$$

$$-\delta p_{t'} \leq 0 \quad (9.13b)$$

$$-\delta p_{t'} \leq -(p_{t'} - p_h). \quad (9.13c)$$

The last constraints define symmetric reserve $\forall t' \in \mathcal{T}_a^m(t)$

$$r_{d,t'}^{s+} \leq \frac{(s_{d,t'} - \underline{S}_d) \eta_d^{\text{dis}}}{\Delta\tau} \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.14a)$$

$$r_{d,t'}^{s+} \leq \underline{P}_d (1 - a_{d,t'}^{\text{dis}}) \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.14b)$$

$$r_{d,t'}^{s-} \leq \frac{(\bar{S}_d - s_{d,t'})}{\eta_d^{\text{cha}} \Delta\tau} \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.14c)$$

$$r_{d,t'}^{s-} \leq \bar{P}_d (1 - a_{d,t'}^{\text{cha}}) \quad \forall d \in \mathcal{D}^{\text{sto}} \quad (9.14d)$$

$$r^{\text{sym, OP}} \leq \sum_{d \in \mathcal{D}^{\text{sto}}} r_{d,t'}^{s+} + \sum_{d \in \mathcal{D}^{\text{ste}}} P_{d,t'}^{\text{ste}} (1 - a_{d,t'}^{\text{ste}}) + \sum_{d \in \mathcal{D}^{\text{nst}}} P_{d,t'}^{\text{nst}} (1 - a^{\text{nst}}) + \sum_{d \in \mathcal{D}^{\text{she}}} C_{d,t'}^{\text{she}} (1 - a_{d,t'}^{\text{she}}) \quad (9.14e)$$

$$r^{\text{sym, OP}} \leq \sum_{d \in \mathcal{D}^{\text{sto}}} r_{d,t'}^{s-} + \sum_{d \in \mathcal{D}^{\text{ste}}} P_{d,t'}^{\text{ste}} a_{d,t'}^{\text{ste}} + \sum_{d \in \mathcal{D}^{\text{nst}}} P_{d,t'}^{\text{nst}} a^{\text{nst}} + \sum_{d \in \mathcal{D}^{\text{she}}} C_{d,t'}^{\text{she}} a_{d,t'}^{\text{she}}. \quad (9.14f)$$

9.3.4 RTP formulation

The RTP objective function implemented for the case study is

$$J_t^{\text{RTO}} = C_t^{\text{RTO}} + D_t^{\text{RTO}} + v_{\tau(t)}(s_{\tau(t)}) \quad (9.15)$$

with Real-Time Optimizer (RTO) the name of this controller. C_t^{RTO} models the immediate costs, D_t^{RTO} the delayed costs and $v_{\tau(t)}(s_{\tau(t)})$ the cost-to-go function of the state of the system at time t within a current market period. C_t^{RTO} is the same as $C_{t'}^{\text{OP}}$ by replacing t' by t , $\Delta\tau$ by Δt and considering only one period of time t . D_t^{RTO} is composed of the peak cost and symmetric reserve penalty costs

$$D_t^{\text{RTO}} = \pi^{\text{p}} \delta p_{\tau(t-\Delta\tau), \tau(t)} + s_t^{\text{TSO}} \pi_{\text{RTO}}^{\text{s}} \Delta r^{\text{sym}}, \quad (9.16)$$

$\delta p_{\tau(t-\Delta\tau), \tau(t)}$ is the peak difference between the previous maximum historic peak p_h and the current peak within the market period computed by RTO. The difference with OP relies on its computation as at t the market period is not finished. Thus, the peak within this market period is made of two parts. (1) The peak from the beginning of the market period to t . (2) The peak from the actions taken from t to the end of the market period. Δr^{sym} is the difference between the symmetric reserve computed by OP and the current reserve within the market period computed by RTO. s_t^{TSO} is the reserve activation signal to activate the tertiary symmetric reserve. It is set by the TSO, 0 if activated, else 1. The activation occurs at the beginning of the following market period.

9.3.5 RTO constraints

The set of constraints that defines the bounds on state and action variables and the energy flows are the same as the OP (9.10) and (9.11) by replacing t' by t , $\Delta\tau$ by Δt and considering only one period of time t . The following constraint describes the dynamics of the state of charge $\forall d \in \mathcal{D}^{sto}$ and $\forall t \in \mathcal{T}_i(t)$

$$s_{d,\tau(t)} - \Delta t (\bar{P}_d \eta_d^{\text{cha}} a_{d,t}^{\text{cha}} - \frac{P_d}{\eta_d^{\text{dis}}} a_{d,t}^{\text{dis}}) = S_{d,t}^{\text{init}}. \quad (9.17)$$

The set of constraints related to the peak power $\forall t \in \mathcal{T}_i(t)$

$$(i_t^{\text{gri}} - e_t^{\text{gri}}) / \Delta t \leq p_{t,\tau(t)} \quad (9.18a)$$

$$-\delta p_{\tau(t)} \leq 0 \quad (9.18b)$$

$$-\delta p_{\tau(t)} \leq -(p_{\tau(t-\Delta\tau),\tau(t)} - p_h) \quad (9.18c)$$

$$p_{\tau(t-\Delta\tau),\tau(t)} = \beta p_{\tau(t-\Delta\tau),t} + (1 - \beta) p_{t,\tau(t)} \quad (9.18d)$$

with $\beta = 1 - \Delta t / \Delta\tau$. The last set of constraints defining the symmetric reserve are the same as the OP (9.14) by replacing t' by t , $r^{\text{sym, OP}}$ by $r^{\text{sym, RTO}}$ and adding $\forall t \in \mathcal{T}_i(t)$

$$-\Delta r^{\text{sym}} \leq 0 \quad (9.19a)$$

$$-\Delta r^{\text{sym}} \leq -(r^{\text{sym, OP}} - r^{\text{sym, RTO}}). \quad (9.19b)$$

9.4 Test description

Our case study is the MiRIS microgrid located at the John Cockerill Group's international headquarters in Seraing, Belgium³. It is composed of PV, several energy storage devices, and a non-sheddable load. The load and PV data we use come from on-site monitoring. All data, including the weather forecasts, are available on the Kaggle platform⁴. The case study consists of comparing RTO-OP to a Rule-Based Controller (RBC) for three configurations of the installed PV capacity, cf. Table 9.1. The RBC prioritizes the use of PV production for the supply of electrical demand. If the microgrid is facing a long position, it charges the battery. Furthermore, if this one is fully charged, it exports to the main grid. If the microgrid is facing a short position, it prioritizes using the battery to supply the demand. Moreover, if this one

³<https://johncockerill.com/fr/energy/stockage-denergie/>


⁴<https://www.kaggle.com/jonathandumas/liege-microgrid-open-data>

is fully discharged, it imports from the main grid. This controller does not take into account any future information, *e.g.*, PV, consumption forecasts, energy prices, or market information such as the peak of the symmetric reserve. Case 3 results from a sizing study that defined the optimal device sizes given the PV and consumption data. The sizing methodology used is described in Dakir and Cornélusse [38].

Figure 9.3 shows the PV & consumption data over the simulation period: from *May 20, 2019* to *June 16, 2019*. The selling price π^e is constant, and the purchasing price is composed of a day π_d^i and night prices π_n^i . Day prices apply from 5 a.m. to 8 p.m. (UTC) during the weekdays, and night prices apply from 8 p.m. to 5 a.m. during weekdays and the entire weekend. The peak mechanism is taken into account with a constant peak price π^p and an initial maximum historic peak p_h . Storage systems are initially fully charged. The PV and consumption data have a 1-second resolution, meaning the RTO could compute its optimization problem each five to ten seconds in operational mode. CPLEX 12.9 is used to solve all the optimization problems on an Intel Core i7-8700 3.20 GHz-based computer with 12 threads and 32 GB of RAM. The average computation time per optimization problem composed of the OP and RTO is a few seconds. However, to maintain a reasonable simulation time, RTO is called every minute. The dataset is composed of 28 days with an average computation time of two hours to solve 1440 optimization problems per day, with a one-minute resolution, leading to two days for the entire dataset. The OP computes a quarterly planning corresponding to the Belgian market period. The computation time of the RTO on a regular computer is around a few seconds and the OP around twenty seconds. In total, the simulation computation time is up to a few hours. The OP computes quarterly planning based on PV and consumption twenty-four ahead forecasts. The weather-based forecast methodology is described in detail in Section 3.3 of Chapter 3. Two "classic" deterministic techniques are implemented, a Recurrent Neural Network (RNN) with the Keras Python library [33] and a Gradient Boosting Regression (GBR) with the Scikit-learn Python library [140]. These models use as input the weather forecasts provided by the Laboratory of Climatology of the Liège University, based on the MAR regional climate model [59]. It is an atmosphere model designed for meteorological and climatic research, used for a wide range of applications, from km-scale process studies to continental-scale multi-decade simulations. To estimate the impact of the PV and consumption forecast errors on the controllers, the simulation is performed with the OP having access to the PV and consumption future values (RTO-OP*). Then, the simulation is performed with the symmetric reserve mechanisms to cope with the

Case	PV_p	\overline{PV}	PV_{max}	PV_{min}	PV_{std}
1	400	61	256	0	72
2	875	133	561	0	157
3	1750	267	1122	0	314
Case	C_p	\overline{C}	C_{max}	C_{min}	C_{std}
1 - 3	1000	153	390	68	72
Case	S_p	$\overline{S}, \underline{S}$	$\underline{P}, \overline{P}$	η^{cha}, η^{dis}	S^{init}
1 - 3	1350	1350, 0	1350, 1350	0.95, 0.95	100
Case	p_h, π^P	I^{cap}	E^{cap}	π_d^i, π_n^i	π^e
1 - 3	150, 40	1500	1500	0.2, 0.12	0.035

Table 9.1 Case studies parameters and data statistics.

Fig. 9.3 Top: PV & consumption simulation data. Bottom: zoom on *June 12, 2019*.

forecast errors. A constant symmetric reserve price π_{OP}^s for the OP and a penalty reserve π_{RTO}^s for the RTO are set to 20 (€/ kW).

9.5 Numerical results

9.5.1 No symmetric reserve

Table 9.2 provides the simulation results without taking into account the symmetric reserve. The smaller the PV installed capacity, the higher the peak and energy costs. The RTO-OP* provides the minimal peak cost, whereas the RBC provides the minimal energy cost on all cases. However, RTO-OP* achieves the minimal total cost, composed

Case 1	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RBC	10.13	6.68	16.81	167	61	0
RTO-OP ^{RNN}	10.37	3.62	13.99	91	64	1
RTO-OP ^{GBR}	10.25	5.27	15.53	132	63	1
RTO-OP*	10.24	0.99	11.23	25	64	1
Case 2	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RBC	3.19	4.85	8.04	121	22	7
RTO-OP ^{RNN}	4.78	2.87	7.65	72	31	15
RTO-OP ^{GBR}	4.30	4.90	9.2	123	28	13
RTO-OP*	4.06	0	4.06	0	26	10
Case 3	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RBC	-2.13	4.12	1.99	105	3	77
RTO-OP ^{RNN}	-1.66	4.12	2.46	105	7	80
RTO-OP ^{GBR}	-1.67	4.23	2.56	106	7	81
RTO-OP*	-1.90	0	0	0	5	79

Table 9.2 Results without symmetric reserve.

of the energy and peak costs. This simulation illustrates the impact of the forecasts on the RTO-OP behavior. The RNN forecaster provides the best results, but the RTO-OP^{RNN} is still a long way to manage the peak as RTO-OP* due to the forecasting errors. The peak cost strongly penalizes the benefits as it applies to the entire year ahead once it has been reached.

In case 3, all the controllers except RTO-OP* reached the maximum peak on *June 12*, 2019 around 10:30 a.m. as shown on Figure 9.4. Figure 9.3 shows a sudden drop in the PV production around 10 a.m. that is not accurately forecasted by the RNN and GBR forecasters as shown in Figure 3.2. This prediction leads to the non-accurate planning of OP. Thus, the RTO cannot anticipate this drop and has to import energy to balance the microgrid at the last minute. Figure 9.5 shows the controllers behavior on *June 12*, 2019 where the peak is reached. In case 2, all controllers reached the same peak as in case 3 except RTO-OP^{RNN} that reached a smaller one on *June 5*, 2019. The forecast's accuracy explains this behavior as in case 3. Finally, in case 1, each controller reached a different peak. The smallest one is achieved by the RTO-OP*, followed by the RTO-OP^{RNN}. These cases show that the RTO-OP controller optimizes

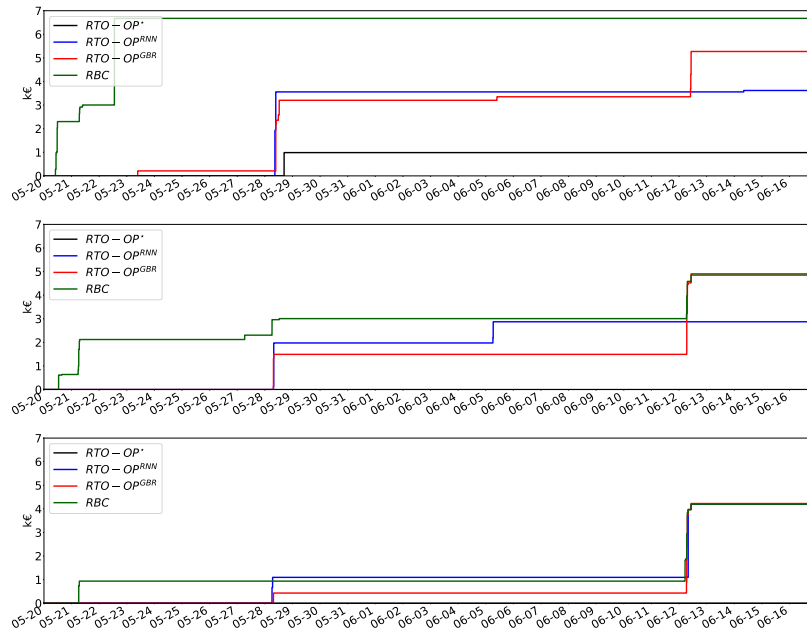


Fig. 9.4 Case 1 (top), 2 (middle), 3 (bottom) cumulative peak costs.

PV-storage usage and thus requires less installed PV capacity for a given demand level. This result was expected as the peak management is not achieved by the RBC and becomes critical when the PV production is smaller than the consumption. This simulation also demonstrates the forecast accuracy impact on the RTO-OP behavior.

9.5.2 Results with symmetric reserve

Table 9.3 provides the simulation results by taking into account the symmetric reserve. Figure 9.6 depicts on case 3 the behavior differences between RTO-OP^{RNN} without and with symmetric reserve. Figures 9.7 and 9.8 show the SOC and peaks costs evolution of case 2 & 1. The controller tends to maintain a storage level that allows RTO-OP^{RNN} to better cope with forecast errors. Indeed for case 3, there is no more peak reached by RTO-OP^{RNN} , only 1 kW for case 2, and it has been almost divided by two for case 1. However, this behavior tends to increase the energy cost if the PV production is large compared to the consumption, such as in case 3. Indeed, the controller will tend to store more energy in the battery instead of exporting it. RTO-OP^* did not perform better with the symmetric reserve. The symmetric reserve competes with the peak management, and the RTO-OP^* tends not to discharge the battery entirely even if it is required to avoid a peak. In case 2, the peak is reached on *June 12, 2019* around 08:00. The controller could have avoided it by totally discharging the battery but did

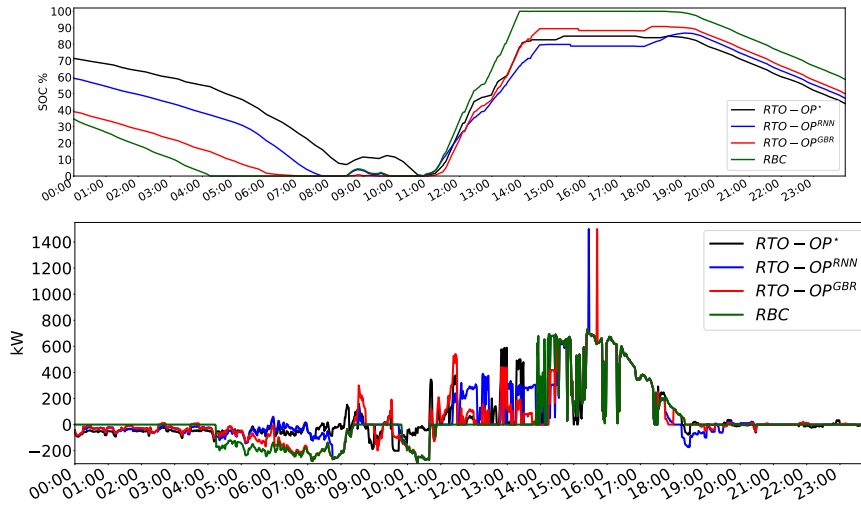


Fig. 9.5 Case 3 SOC (top) and net export power (bottom) on *June 12, 2019*.




Fig. 9.6 Case 3 SOC comparison for RTO-OP^{RNN} with and without symmetric reserve.

not maintain the reserve level. It is the same behavior in case 1, where the peak could have been limited if all the battery was discharged. There is an economic trade-off to manage the peak and the reserve simultaneously. It depends on the valorization or not on the market of the symmetric reserve. The reserve can also be valorized internally to cope with non or complex forecastable events. Such as a sudden drop in export or import limits due to loss of equipment or grid congestion.

9.6 Conclusions

A two-level value function-based approach is proposed as a solution method for a multi-resolution microgrid optimization problem. The value function computed by the operational planner based on PV and consumption forecasts allows coping with the forecasting uncertainties. The real-time controller solves an entire optimization problem, including the future information propagated by the value function. This approach is tested on the MiRIS microgrid case study with PV and consumption data

Case 1	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RTO-OP ^{RNN}	10.50	2.12	12.62	53	65	3
RTO-OP*	10.47	2.75	13.22	69	65	2
Case 2	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RTO-OP ^{RNN}	5.33	0.04	5.37	1	41	27
RTO-OP*	4.78	0.99	5.77	25	35	20
Case 3	c_E	c_p	c_t	Δ_p	I_{tot}	E_{tot}
RTO-OP ^{RNN}	-0.04	0	-0.04	0	24	99
RTO-OP*	-0.15	0	-0.15	0	23.2	98

Table 9.3 Results with symmetric reserve.

monitored on-site. The results demonstrate the efficiency of this method to manage the peak in comparison with a Rule-Based Controller. This test case is completely reproducible as all the data used are open, PV, consumption monitored, and forecasted, including the weather forecasts. The proposed method can be extended in three ways. First, a stochastic formulation of the operational planning problem to cope with probabilistic forecasts. Second, adding the balancing market mechanisms. Finally, considering a community composed of several entities inside the microgrid.




Fig. 9.7 Case 1 (top) and 2 (bottom) cumulative peak costs.




Fig. 9.8 Case 1 (top) and 2 (bottom) SOC.

Chapter 10

Capacity firming using a stochastic approach

Overview

This Chapter proposes a stochastic approach to address the energy management of a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. Both deterministic and stochastic approaches result in optimization problems formulated as quadratic problems with linear constraints. The considered case study is a real microgrid with PV production monitored on-site.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Bertrand Cornélusse, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Stochastic and deterministic formulations for capacity firming nominations. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–7. IEEE, 2020. URL <https://arxiv.org/abs/2106.02425>.

“*We are our choices.*”

— Jean-Paul Sartre

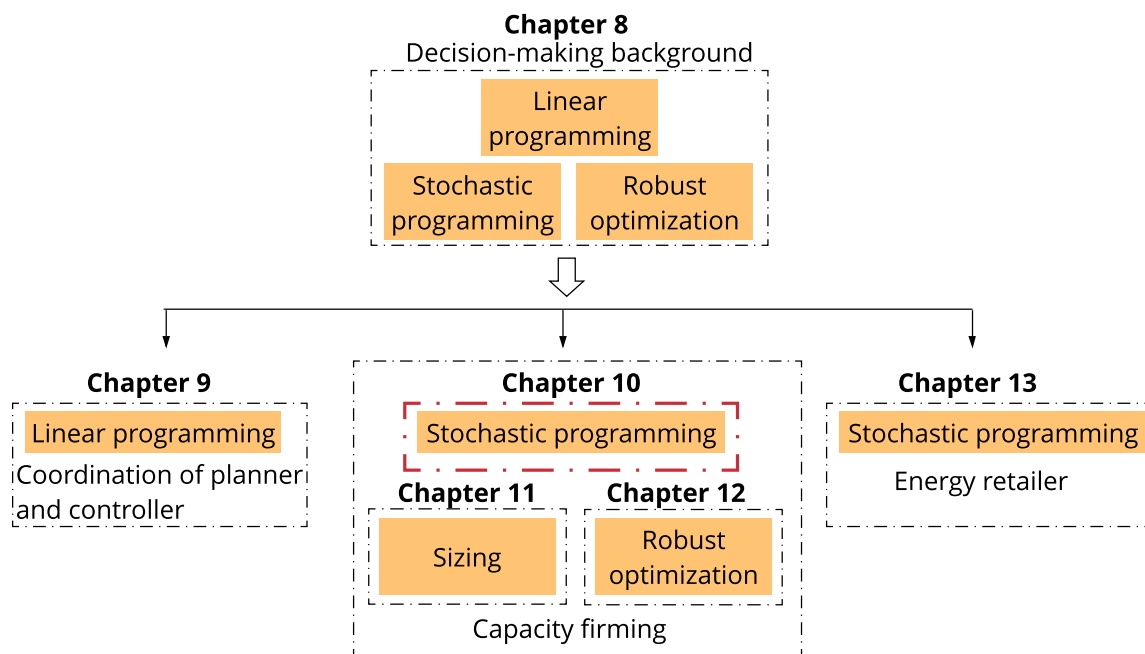


Fig. 10.1 Chapter 10 position in Part II.

The capacity firming framework is mainly designed for isolated markets, such as the Overseas France islands. For instance, the French Energy Regulatory Commission (CRE) publishes capacity firming tenders and specifications. The system considered is a grid-connected renewable energy power plant, *e.g.*, photovoltaic or wind-based, with a battery energy storage system (BESS) for firming the renewable generation. At the tendering stage, offers are selected on the electricity selling price. Then, the successful tenderer builds its plant and sells the electricity exported to the grid at the contracted selling price, but according to a well-defined daily engagement and penalization scheme specified in the tender specifications. The electricity injected in or withdrawn from the grid must be nominated the day-ahead, and engagements must satisfy ramping power constraints. The remuneration is calculated a posteriori by multiplying the realized exports by the contracted selling price minus a penalty. The deviations of the realized exports from the engagements are penalized through a function specified in the tender. A peak option can be activated in the contract for a significant selling price increase during a short period defined a priori. Therefore, the BESS must shift the renewable generation during peak hours to maximize revenue and manage renewable energy uncertainty.

The problem of modeling a two-phase engagement/control with an approach dealing with uncertainty in the context of the CRE capacity framework is still an open issue. This framework has received less attention in the literature than more traditional energy markets such as day-ahead and intraday markets of European countries.

The optimal day-ahead bidding strategies of a plant composed of only a production device have been addressed in, *e.g.*, Pinson et al. [143], Bitar et al. [23], Giannitrapani et al. [65, 66]. The optimal offer turns out to be a suitable percentile of the PV/wind power cumulative distribution function. Under the assumption of time-invariant power generation statistics, the cumulative distribution functions can be estimated from historical data of the power generated by the plant. This assumption is not always justified, especially for PV power generation. In Giannitrapani et al. [66], the authors investigate two approaches to properly take into account the effects of seasonal variation and non-stationary nature of PV power generation in the estimation of PV power statistics. However, incorporating energy storage in the framework is still an open problem, and the literature provides several approaches and methodologies to this end. An optimal power management mechanism for a grid-connected PV system with storage is implemented in [150] using Dynamic Programming (DP) and is compared with simple ruled-based management. The sizing and control of an energy storage system to mitigate wind power uncertainty is addressed by Haessig [80], Haessig et al.

[81, 82] using stochastic dynamic programming (SDP). The framework is similar to the CRE PV capacity firming tender with a wind farm operator committed on a day-ahead basis to a production engagement. Finally, three distinct optimization strategies, mixed-integer quadratic programming, simulation-based genetic algorithm, and expert-based heuristic, are empirically compared by N’Goran et al. [128] in the CRE framework.

This study addresses the energy management of a grid-connected PV plant and BESS. This topic is studied within the capacity firming specifications of the CRE, in line with the tender AO-CRE-ZNI 2019 published on *July 12, 2019*, using the MiRIS microgrid case study. The capacity firming problem can be decomposed into two steps. The first step consists of computing the day-ahead nominations. The second step consists of computing the renominations and the set-points in real-time to minimize the energy and ramp power deviations from nominations. This study focuses on the first step and proposes both a stochastic and a deterministic formulation. The main goal of this study is to validate the stochastic approach by using an ideal predictor providing unbiased PV scenarios. Thus, the BESS efficiencies are perfect, and the degradation is not taken into account for the sake of simplicity. Different levels of prediction accuracy are evaluated. Then, the results are compared with those of the deterministic formulation, assuming perfect forecasts returned by an oracle. Both deterministic and stochastic approaches result in optimization problems formulated as quadratic problems with linear constraints. The considered case study is a real microgrid with PV production monitored on-site.

The study is organized as follows. Section 10.1 provides the notation that is also used for Chapters 11 and 12. Section 10.2 describes the capacity firming framework. Section 10.3 proposes the deterministic and stochastic formulations of the nomination process. Section 10.4 introduces the MiRIS microgrid case study and presents the results. Conclusions are drawn in Section 10.5. Appendix 10.6 describes the methodology to generate the set of unbiased PV scenarios.

10.1 Notation

10.1.1 Sets and indices

Name	Description
------	-------------

t	Time period index.
ω	PV scenario index.
q	PV quantile index.
T	Number of time periods per day.
\mathcal{T}	Set of time periods, $\mathcal{T} = \{1, 2, \dots, T\}$.
\mathcal{D}	$\cup_{i=1}^{i=D} \mathcal{T}_i$ set of D days of market periods.
Ω	PV scenarios set, $\Omega = \{1, 2, \dots, \#\Omega\}$.
$\#\Omega$	Number of PV scenarios in Ω .
\mathcal{U}	PV uncertainty set for robust optimization.

10.1.2 Parameters

Name	Description	Unit
P_c	PV installed capacity.	kWp
\bar{S}	BESS installed capacity.	kWh
y_t^m	Measured power at the grid connection point.	kW
$y_t^{\text{PV,m}}$	Measured PV generation.	kW
\hat{y}_t^{PV}	PV point forecast.	kW
$\hat{y}_{t,\omega}^{\text{PV}}$	PV scenario ω .	kW
$\hat{y}_t^{\text{PV},(q)}$	PV quantile forecast q .	kW
α_ω	Probability of PV scenario ω	-
$y_{\text{max}}^{\text{dis}}, y_{\text{max}}^{\text{cha}}$	BESS maximum (dis)charging power.	kW
$\eta^{\text{dis}}, \eta^{\text{cha}}$	BESS (dis)charging efficiency.	-
$s^{\text{min}}, s^{\text{max}}$	BESS minimum/maximum capacity with $s^{\text{max}} \leq \bar{S}$.	kWh
$s^{\text{ini}}, s^{\text{end}}$	BESS initial/final state of charge.	kWh
$x_t^{\text{min}}, x_t^{\text{max}}$	Minimum/maximum engagement.	kW
pP_c	Engagement tolerance, $0 \leq p \leq 1$, kW.	
ΔX	Ramping-up and down limits for the engagement.	kW
$y_t^{\text{min}}, y_t^{\text{max}}$	Minimum/maximum power at the grid connection point.	kW
π_t	Contracted selling price.	€/kWh
π^e	Slack price.	$\frac{\text{€}}{\text{kWh}^2}$
$\pi_{\bar{S}}$	BESS CAPEX price.	€/kWh
Δt	Time period duration.	minutes
c	Penalty function.	€
$u_t^{\text{min}}, u_t^{\text{max}}$	Minimum/maximum of the PV uncertainty interval.	kW
Γ	Uncertainty budget for robust optimization.	-

d_q, d_Γ	Uncertainty and budget depths for robust optimization.	%
M_t^-, M_t^+	Big-M's values for robust optimization.	-

10.1.3 Variables

For the sake of clarity the subscript ω is omitted.

Name	Range	Description	Unit
x_t	$[x_t^{\min}, x_t^{\max}]$	Engagement.	kW
y_t	$[y_t^{\min}, y_t^{\max}]$	Net power at the grid connection point.	kW
y_t^{PV}	$[0, \hat{y}_t^{\text{PV}}]$	PV generation.	kW
y_t^{cha}	$[0, y_{\max}^{\text{cha}}]$	BESS charging power.	kW
y_t^{dis}	$[0, y_{\max}^{\text{dis}}]$	BESS discharging power.	kW
y_t^s	$[s^{\min}, s^{\max}]$	BESS state of charge.	kWh
d_t^-, d_t^+	\mathbb{R}_+	Short/long deviation.	kW
y_t^b	$\{0, 1\}$	BESS binary variable.	-
z_t	$\{0, 1\}$	PV uncertainty set binary variable for robust optimization.	-
α_t	$[M_t^-, M_t^+]$	Variables to linearize $z_t \phi_t^{y^{\text{PV}}}$ for robust optimization.	-

Dual variables, and corresponding constraints

Dual variables of constraints are indicated with brackets $[\cdot]$.

Name	Range	Description
$\phi_t^{\text{cha}}, \phi_t^{\text{dis}}$	\mathbb{R}^-	Maximum storage (dis)charging.
$\phi_t^{s^{\min}}, \phi_t^{s^{\max}}$	\mathbb{R}^-	Minimum/maximum storage capacity.
ϕ_t^y	\mathbb{R}	Net power balance.
$\phi_t^{y^{\min}}, \phi_t^{y^{\max}}$	\mathbb{R}^-	Minimum/maximum net power.
$\phi_t^{s^{\text{ini}}}, \phi_t^{s^{\text{end}}}$	\mathbb{R}^-	Initial/final state of charge.
ϕ_t^s	\mathbb{R}^-	BESS dynamics.
$\phi_t^{d^-}, \phi_t^{d^+}$	\mathbb{R}^-	Under/overproduction.
$\phi_t^{y^{\text{PV}}}$	\mathbb{R}^-	Renewable generation.

10.2 The Capacity Firming Framework

The capacity firming framework can be decomposed into a day-ahead engagement process, Section 10.2.1, and a real-time control process, Section 10.2.2. Each day is discretized in T periods of duration Δt . In the sequel, the period duration is the same for day-ahead engagement and the real-time control, t is used as a period index, and \mathcal{T} is the set of periods in a day.




Fig. 10.2 Day-ahead nomination process.

10.2.1 Day-ahead engagement

Each day, the operator of the renewable generation plant is asked to provide the generation profile to be followed the next day to the grid operator, based on renewable generation forecasts. More formally, a planner computes on a day-ahead basis, before a deadline, a vector of engagements composed of T values $[x_1, \dots, x_T]^T$. Figure 10.2 illustrates the day-ahead nomination process. The grid operator accepts the engagements if they satisfy the constraints

$$|x_t - x_{t-1}| \leq \Delta X, \quad \forall t \in \mathcal{T} \setminus \{1\} \quad (10.1a)$$

$$-x_t \leq -x_t^{\min}, \quad \forall t \in \mathcal{T} \quad (10.1b)$$

$$x_t \leq x_t^{\max}, \quad \forall t \in \mathcal{T}, \quad (10.1c)$$

with ΔX a ramping power constraint, a fraction of the total installed capacity P_c determined at the tendering stage and imposed by the grid operator.

10.2.2 Real-time control

Then, in real-time, a receding-horizon controller computes at each period the generation level and the charge or discharge set-points from t to T , based on forecasts of renewable generation and the engagements. Only the set-points of the first period are applied to the system. The remuneration is calculated ex-post based on the realized power y_t^m at the grid coupling point. For a given control period, the net remuneration r_t of the plant is the gross revenue $\Delta t \pi_t y_t^m$ minus a penalty $c(x_t, y_t^m)$, with π_t the contracted selling price set at the tendering stage

$$r_t = \Delta t \pi_t y_t^m - c(x_t, y_t^m), \quad \forall t \in \mathcal{T}. \quad (10.2)$$

The penalty function c depends on the specifications of the tender. For the sake of simplicity for the rest of this Chapter, c is assumed to be symmetric, convex, and

quadratic piecewise-linear

$$c(x_t, y_t^m) = \pi^e (\Delta t)^2 \left(\max \left(0, |x_t - y_t^m| - pP_c \right) \right)^2 \quad (10.3)$$

with $0 \leq p \leq 1$, and π^e is a slack price ($\text{€}/\text{kWh}^2$).

10.3 Problem formulation

The problem statement follows the abstract formulation defined in Chapter 9 where a global microgrid control problem can be defined as operating a microgrid safely and in an economically efficient manner. In the capacity firming context, a two-stage approach is considered with a planner and a controller. A quadratic formulation with linear constraints models the CRE non-convex penalty. In this study, the planner and controller optimize economic criteria, which are only related to active power. The ancillary or grid services are not in the scope of the capacity firming specifications. The BESS degradation is not taken into account. The planner and controller horizons are cropped to twenty-four hours.

Deterministic (D) and stochastic (S) formulations of the day-ahead nomination problem are compared. The deterministic formulation is used as a reference to validate the stochastic approach by considering perfect knowledge of the future (D^*). In this Chapter, both approaches consider only exports to the grid¹. The optimization variables and the parameters are defined in Section 10.1.

10.3.1 Deterministic approach

The objective function J_D to minimize is the opposite of the net revenue

$$J_D(x_t, y_t) = \sum_{t \in \mathcal{T}} -\pi_t \Delta t y_t + \pi^e (\Delta t)^2 [(d_t^-)^2 + (d_t^+)^2]. \quad (10.4)$$

¹The imports from the grid are allowed only under specific conditions into the contract.

The deterministic formulation is the following Mixed-Integer Quadratic Program (MIQP)

$$\begin{aligned} \min_{x_t \in \mathcal{X}, y_t \in \Omega(x_t, \hat{y}_t^{\text{PV}})} J_D(x_t, y_t) \\ \mathcal{X} = \{x_t : (10.6)\} \\ \Omega(x_t, \hat{y}_t^{\text{PV}}) = \{y_t : (10.7) - (10.11)\}, \end{aligned} \quad (10.5)$$

where \mathcal{X} and $\Omega(x_t, \hat{y}_t^{\text{PV}})$ are the sets of feasible engagements x_t and dispatch solutions y_t for a fixed engagement x_t and renewable generation point forecast \hat{y}_t^{PV} . The optimization variables of (10.5) are the engagement variables x_t , the dispatch variables y_t (the net power at the grid connection point), y_t^{dis} (BESS discharging power), y_t^{cha} (BESS charging power), s_t (BESS state of charge), y_t^b (BESS binary variables), y_t^{PV} (renewable generation), and d_t^-, d_t^+ (deviation variables) (cf. the notation Section 10.1). From (10.1), the engagement constraints are

$$x_t - x_{t-1} \leq \Delta X, \quad \forall t \in \mathcal{T} \setminus \{1\} \quad (10.6a)$$

$$x_{t-1} - x_t \leq \Delta X, \quad \forall t \in \mathcal{T} \setminus \{1\} \quad (10.6b)$$

$$-x_t \leq -x_t^{\min}, \quad \forall t \in \mathcal{T} \quad (10.6c)$$

$$x_t \leq x_t^{\max}, \quad \forall t \in \mathcal{T}. \quad (10.6d)$$

The ramping constraint on x_1 is deactivated to decouple consecutive days of simulation. In reality, the updated value of the last engagement of the previous day would be taken to satisfy the constraint. The set of constraints that bound y_t^{cha} , y_t^{dis} , and y_t^s variables are $\forall t \in \mathcal{T}$

$$y_t^{\text{cha}} \leq y_t^b y_{\max}^{\text{cha}} \quad [\phi_t^{\text{cha}}] \quad (10.7a)$$

$$y_t^{\text{dis}} \leq (1 - y_t^b) y_{\max}^{\text{dis}} \quad [\phi_t^{\text{dis}}] \quad (10.7b)$$

$$-s_t \leq -s^{\min} \quad [\phi_t^{s^{\min}}] \quad (10.7c)$$

$$s_t \leq s^{\max}, \quad [\phi_t^{s^{\max}}] \quad (10.7d)$$

where y_t^b are binary variables that prevent the simultaneous charge and discharge of the BESS. The power balance equation and the constraints on the net power at the

grid connection point are $\forall t \in \mathcal{T}$

$$y_t = y_t^{\text{pv}} + (y_t^{\text{dis}} - y_t^{\text{cha}}) \quad [\phi_t^y] \quad (10.8a)$$

$$-y_t \leq -y^{\text{min}} \quad [\phi_t^{y^{\text{min}}}] \quad (10.8b)$$

$$y_t \leq y^{\text{max}}. \quad [\phi_t^{y^{\text{max}}}] \quad (10.8c)$$

The dynamics of the BESS state of charge are

$$s_1 - s^{\text{ini}} = \Delta t (\eta^{\text{cha}} y_1^{\text{cha}} - \frac{y_1^{\text{dis}}}{\eta^{\text{dis}}}) \quad [\phi^{s^{\text{ini}}}] \quad (10.9a)$$

$$s_t - s_{t-1} = \Delta t (\eta^{\text{cha}} y_t^{\text{cha}} - \frac{y_t^{\text{dis}}}{\eta^{\text{dis}}}), \quad \forall t \in \mathcal{T} \setminus \{1\} \quad [\phi_t^s] \quad (10.9b)$$

$$s_T = s^{\text{end}} = s^{\text{ini}}, \quad [\phi^{s^{\text{end}}}] \quad (10.9c)$$

where the parameters s^{end} and s^{ini} are introduced to decouple consecutive days of simulation. In reality, s^{ini} would be the updated value of the last measured state of charge of the previous day. The variables d_t^-, d_t^+ are defined $\forall t \in \mathcal{T}$ to model the penalty

$$-d_t^- \leq (y_t - (x_t - pP_c)) \quad [\phi_t^{d^-}] \quad (10.10a)$$

$$-d_t^+ \leq ((x_t + pP_c) - y_t), \quad [\phi_t^{d^+}] \quad (10.10b)$$

with $0 \leq p \leq 1$. Finally, the PV generation is bounded $\forall t \in \mathcal{T}$ by the point forecast \hat{y}_t^{pv}

$$y_t^{\text{pv}} \leq \hat{y}_t^{\text{pv}}. \quad [\phi_t^{y^{\text{pv}}}] \quad (10.11)$$

10.3.2 Deterministic approach with perfect forecasts

With perfect forecasts, (10.5) becomes

$$\min_{x_t \in \mathcal{X}, y_t \in \Omega(x_t, \hat{y}_t^{\text{pv}} = y_t^{\text{pv}, \text{m}})} J_{D^*}(x_t, y_t), \quad (10.12)$$

with $\hat{y}_t^{\text{pv}} = y_t^{\text{pv}, \text{m}} \forall t \in \mathcal{T}$ in (10.11).

10.3.3 Stochastic approach

In the stochastic formulation, the objective is given by

$$J_S = \mathbb{E} \left[\sum_{t \in \mathcal{T}} -\pi_t \Delta t y_t + c(x_t, y_t) \right] \quad (10.13)$$

where the expectation is taken with respect to \hat{y}_t^{pv} . Using a scenario-based approach, (10.13) is approximated by

$$J_S \approx \sum_{\omega \in \Omega} \alpha_\omega \sum_{t \in \mathcal{T}} \left[-\pi_t \Delta t y_{t,\omega} + \pi^e (\Delta t)^2 \left((d_{t,\omega}^-)^2 + (d_{t,\omega}^+)^2 \right) \right], \quad (10.14)$$

with α_ω the probability of scenario $\omega \in \Omega$, and $\sum_{\omega \in \Omega} \alpha_\omega = 1$. Then, the problem to solve becomes

$$\min_{x_t \in \mathcal{X}, y_{t,\omega} \in \Omega(x_t, \hat{y}_{t,\omega}^{\text{pv}})} J_S(x_t, y_{t,\omega}). \quad (10.15)$$

All the optimization variables but x_t are now defined $\forall \omega \in \Omega$.

10.3.4 Evaluation methodology

The second step of capacity firming, *i.e.*, computing the set-points in real-time, is required to assess the quality of the nomination process. However, since this study focuses on the computation of day-ahead engagements, we simulate the second step with an ideal real-time controller² once the engagements are fixed. The methodology to assess the engagements consists of solving

$$J^{\text{eval}} = \sum_{t \in \mathcal{T}} -\pi_t x_t + c(x_t, y_t) \quad (10.16)$$

s.t (10.7)-(10.11) with $\hat{y}_t^{\text{pv}} = y_t^{\text{pv,m}}$ in (10.11) and given engagements x_t previously computed by the planner S. The optimization variables of (10.16) are y_t , d_t^- , d_t^+ , y_t^{pv} , y_t^{cha} , y_t^{dis} , and s_t . The optimal value of J_S^{eval} is compared with the optimal value of J_{D^*} in (10.12).

²Using a real-time controller with intraday forecasts is required to assess the planner-controller. However, this study focus only on the nomination step.




Fig. 10.3 MiRIS February 2019 PV production.

10.4 MiRIS microgrid case study

The MiRIS³ microgrid case study, located at the John Cockerill Group’s international headquarters in Seraing, Belgium, is composed of a PV production plant, a BESS, and a load. For the need of this study, only historical data of PV generation are required. The BESS capacity \bar{S} is 1000 kWh, $s^{\max} = \bar{S}$, and the total PV installed capacity P_c is 2000 kWp. The market period duration Δt is 15 minutes. The simulation dataset \mathcal{D} is the month of February 2019. Figure 10.3 illustrates the MiRIS PV production and Table 10.1 provides some key statistics. Table 10.2 defines the indicators used in this section. It is of paramount importance to notice that the results of this case study

P_c	$\overline{y^{\text{pv,m}}}$	$y_{\text{std}}^{\text{pv,m}}$	$y_{\text{max}}^{\text{pv,m}}$	$y_{\%,\text{max}}^{\text{pv,m}}$	$y_{\text{tot}}^{\text{pv,m}}$
2000	104.6	202.4	988.1	49.4	70.3

Table 10.1 MiRIS February 2019 dataset statistics.

are only valid for this dataset and cannot be extrapolated over an entire year without caution. CPLEX⁴ 12.9 is used to solve all the optimization problems, on an Intel Core i7-8700 3.20 GHz based computer with 12 threads and 32 GB of RAM. Tables 10.3 and 10.4 provide the case study and BESS parameters.

³<https://johncockerill.com/fr/energy/stockage-denergie/>

⁴<https://www.ibm.com/products/ilog-cplex-optimization-studio>

Name	Description	Unit
$\overline{y^{\text{pv},m}}$	Averaged PV generation.	kW
$y_{\text{std}}^{\text{pv},m}$	PV generation standard deviation.	kW
$y_{\text{max}}^{\text{pv},m}$	Maximum PV generation.	kW
$y_{\text{tot}}^{\text{pv},m}$	Total PV energy produced.	MWh
$y_{\%,\text{max}}^{\text{pv},m}$	$y_{\text{max}}^{\text{pv},m}$ divided by the total installed PV capacity P_c .	%
\bar{t}_{CPU}	Averaged computation time per optimization problem.	s
$[X]^D$	Total of a variable X_t : $\sum_{t \in \mathcal{D}} X_t$.	X unit
$[y^{\text{pv}}]^D$	Total PV generation.	MWh
$y_{\%}^{\text{pv}}$	PV generation ratio: $\frac{[y^{\text{pv}}]^D}{[y^{\text{pv},m}]^D}$.	%
$y_{\%}^{\text{cha}}$	BESS charge ratio: $\frac{[y^{\text{cha}}]^D}{[y^{\text{pv}}]^D}$.	%
$y_{\text{max}\%}^{\text{cha}}$	Percentage of days of the dataset where the BESS achieved its maximum storage capacity.	%
$y_{\%}$	Export ratio: $\frac{[y]^D}{[x]^D}$.	%
R_{max}	Maximum achievable revenue: $\pi[y^{\text{pv},m}]^D$.	k€
R^e	Gross revenue: $\pi[y]^D$.	k€
r^e	Maximum achievable revenue ratio: $\frac{R^e}{R_{\text{max}}}$.	%
C^e	Quadratic penalty: $[c]^D$.	k€
$R^{n,e}$	Net revenue with quadratic penalty: $R^e - C^e$.	k€

Table 10.2 Comparison indicators.

π	π^e	Δt	ΔX	x^{\max}	y^{\max}	p
0.045	0.0045	15	10	2000	2000	5%

Table 10.3 Case study parameters.

s^{\max}	s^{\min}	y_{\max}^{cha}	y_{\max}^{dis}	η^{cha}	η^{dis}	s^{ini}	s^{end}
1000	0	1000	1000	1	1	0	0

Table 10.4 BESS parameters.

10.4.1 Results for unbiased PV scenarios with fixed variance

A set of unbiased PV scenarios is generated for several values of the standard deviation σ of the prediction error. Table 10.5 shows the considered values of σ , expressed as a fraction of the actual PV generation. Moreover, Table 10.5 reports the cardinality of the generated scenario sets. Table 10.6 compares the average computation time per optimization problem between planners S and D*. Note, The optimization problem of planner S with $\#\Omega = 1$ has the same number of variables and constraints as the planner D*. The computation time is compatible with a day-ahead process even with 100 scenarios, as it takes on average 7 seconds to compute the nominations for the day-ahead. Table 10.7 and Figure 10.4 provide the results of the ratio indicators, respectively, for the planners D* and S.

For all indicators, the results of both planners are almost equal with the smaller value of σ and the highest value of $\#\Omega$, as expected. On average, the curtailment of PV generation equals 5%. The maximum $y_{\%}$ value is achieved with $\sigma = 3.5\%$ because the nominations are more conservative when the variance increases, leading to a smaller ratio. On average, 30% (27%) of the production, for planner D* (S), is stored in the BESS over the entire dataset. $\bar{S}_{\%}$ is equal to 17.9% (17.9%⁵) for the planner D* (S), meaning the BESS reached its maximum storage level 5 days out of the 28 days of the dataset. In fact, during sunny days, the BESS is fully charged. A larger BESS capacity should decrease the curtailment and improve the gross revenue. Note: it is a winter

⁵The value is the same for the $\#\Omega$ and σ values considered.

σ	3.5%	7%	10.5%	14%
$\#\Omega$	5	10	50	100

Table 10.5 Scenario generation parameters.

$\#\Omega$	1	5	10	50	100
# variables	769	3 457	6 817	33 697	67 297
# constraints	1 248	5 092	9 897	48 337	96 387
$\bar{t}_{\text{CPU S}}$	-	0.3	0.8	3	7
$\bar{t}_{\text{CPU D}^*}$	0.1	-	-	-	-

Table 10.6 Averaged computation times.


$[y^{\text{pv}}]^D$	$y_{\%}^{\text{pv}}$	$y_{\%}^{\text{cha}}$	$y_{\text{max}\%}^{\text{cha}}$	$y_{\%}$
66.7	94.9	29.6	17.9	76.2

Table 10.7 Planner D* ratio indicators.

month where the maximum generated PV power reached only half of the installed PV capacity. During a summer month, the maximum production should reach at least 80% of the total installed capacity on sunny days. Thus, with a storage capacity of 1 MWh, the curtailment is expected to be higher during sunny summer days.


Table 10.8 and Figure 10.5 provide the results of the revenue indicators for the planners D* and S, respectively. It should be noted that in this case, $R^{n,e} = -J_S^{\text{eval}}$. The smallest value of the objective function is achieved by the planner D*. It is followed closely by the planner S, even for the highest value of σ . This result demonstrates the validity of the approach when exploiting an unbiased stochastic predictor.

In terms of net revenue, both planners achieved 93.7% of $R_{\text{max}} \approx 3.16$ k€, which results in a loss of 6.3%. Most of this loss is due to the curtailment of PV generation. For both planners, the net revenue increases with the generation.

Fig. 10.4 Planner S ratio indicator with $\#\Omega = 5, 10, 50, 100$, $\sigma = 3.5, 7, 10.5, 14\%$.

R^e	r^e	C^e	$R^{n,e}$	$J_{D^*}^{\text{eval}}$
3.0	94.9	0.04	2.96	-2.96

Table 10.8 Planner D* revenue indicators.

Fig. 10.5 Planner S revenue indicators with $\#\Omega = 5, 10, 50, 100$, $\sigma = 3.5, 7, 10.5, 14\%$.

For sunny days, the difference between the nominations and the exports is higher than the tolerance just before the production occurs, between 5 and 8 am. Indeed, the planner tends to maximize the revenue by maximizing the exports. However, the ramping power constraints (10.6) impose a maximum difference between two consecutive nominations. To maximize the net revenue over the entire day, the planner computes nominations that are not achievable at the beginning of the day to maximize the exports during the day. This results in a penalty between 5 and 8 am.

10.4.2 BESS capacity sensitivity analysis

The goal is to conduct a sensitivity analysis on the BESS capacity \bar{S} to determine its marginal value and the optimal BESS size \bar{S}^* for a given CAPEX $\pi_{\bar{S}}$. The efficiencies are still assumed to be unitary. s^{ini} and s^{end} are set to 0 kWh. Table 10.9 provides the other BESS parameters for the five cases. The scenarios are generated using $\#\Omega = 100$, and $\sigma = 3.5, 7, 10.5, 14\%$. A new indicator, expressed in k€, is defined to quantify the gain provided by the BESS over fifteen years

$$\Delta R_i^{n,e} = 15 \times 12 \times (R_i^{n,e} - R_5^{n,e}) \quad \forall i \in \{1, 2, 3, 4\}. \quad (10.17)$$


It is a lower bound of the total gain as it relies on the results of a winter month. A summer month should provide higher revenue. Table 10.10 provide the planner D*

Case	\bar{S} [kWh]	s^{\min} [kWh]	y_{\max}^{cha} [kW]	y_{\max}^{dis} [kW]
1	2000	0	2000	2000
2	1000	0	1000	1000
3	500	0	500	500
4	250	0	250	250
5	0	0	0	0

Table 10.9 BESS parameters.

indicators. The results demonstrate the interest in using a BESS to optimize the bidding. The larger the BESS is, the lower the curtailment is. Thus, the net revenue increases with the BESS capacity. The maximum achievable revenue is reached with a storage capacity of 2 MWh. However, the larger the BESS is, the smaller $\Delta R^{n,e}$ increases. It means the marginal benefit decreases with the increase of BESS capacity. A trade-off should be found between the BESS capacity and its CAPEX. Figure 10.6a provides $\Delta R^{n,e}$ and its quadratic interpolation in comparison with two BESS prices $\pi_{\bar{S}} = 0.1$ and 0.228 k€/kWh. The value of the derivative $[\frac{d\Delta R^{n,e}}{d\bar{S}}]_{\bar{S}=0}$ provides the maximum CAPEX that provides a profitable BESS. Then, the optimal storage capacity \bar{S}^* for a given CAPEX is provided solving $[\frac{d\Delta R^{n,e}}{d\bar{S}}]_{\bar{S}} = \pi_{\bar{S}}$. For instance, with a CAPEX of 0.1 k€/kWh, \bar{S}^* is approximately 350 kWh. Figure 10.6b provides the values of $\Delta R^{n,e} - \pi_{\bar{S}}\bar{S}$ with a quadratic interpolation.

Figure 10.7 provides the planner S revenue indicators. The results are still almost identical for all indicators for the smallest value of σ and very close with the highest one, as expected.



(a) Variation of the net revenue.


(b) Optimal BESS.

Fig. 10.6 Optimal BESS for a given CAPEX price.

Case	$[y^{\text{pv}}]^D$	$y_{\%}^{\text{pv}}$	$y_{\%}^{\text{cha}}$	$y_{\text{max}\%}^{\text{cha}}$	$y_{\%}$
1	70.3	100	45.6	17.9	77.7
2	66.7	94.9	29.6	17.9	76.2
3	63.6	90.5	17.3	39.3	76.4
4	61.7	87.7	11.1	46.4	75.2
5	56.0	79.6	-	-	70.4

Case	R^e	C^e	$R^{n,e}$	$J_{D^*}^{\text{eval}}$	$\Delta R^{n,e}$
1	3.16	0.01	3.15	-3.15	128
2	3.0	0.04	2.96	-2.96	94
3	2.86	0.04	2.84	-2.84	72
4	2.77	0.06	2.71	-2.71	49
5	2.52	0.08	2.44	-2.44	0

Table 10.10 Planner D* ratio and revenue indicators BESS capacity sensitivity analysis.

Fig. 10.7 Planner S revenue indicators BESS capacity sensitivity analysis with $\#\Omega = 100$, $\sigma = 3.5, 7, 10.5, 14\%$.

10.5 Conclusions and perspectives

This Chapter addresses the energy management of a grid-connected PV plant coupled with a BESS within the capacity firming framework. The method is composed of two steps: computing the day-ahead nominations, then computing the renominations and the set-points in real-time to minimize the energy and ramp power deviations from nominations. This study investigates the first step by comparing a stochastic and a deterministic formulation. The main goal is to validate the stochastic approach by using an ideal predictor providing unbiased PV scenarios.

The results of the stochastic planner are comparable with those of the deterministic planner, even when the prediction error variance is non-negligible. Finally, the BESS capacity sensitivity analysis results demonstrate the advantage of using a BESS to optimize the bidding day-ahead strategy. However, a trade-off must be found between the marginal gain provided by the BESS and its investment and operational costs.

Several extensions of this work are under investigation. The first is to assess the planner's behavior better using a full year of data. Then, the next challenge is to use a more realistic methodology to generate PV generation scenarios. Several scenario generation approaches could be investigated, based on a point forecast model such as the PVUSA model [46, 19, 20], combined with Gaussian copula [135, 142, 72]. Alternatively, using the deep generative models introduced in Chapter 6. Another challenge is to consider the non-convex penalty function specified by the CRE into the objective. Finally, the last challenge is to investigate the second step of the capacity firming problem, for instance, by adapting the approach implemented in Dumas et al. [54].

10.6 Appendix: PV scenario generation

This Annex describes the methodology to generate the set of unbiased PV scenarios. The goal is to define an ideal unbiased predictor with a fixed variance over all lead times. In this section, let t be the current time index, k be the lead time of the prediction, K be the maximum lead time of the prediction, y_{t+k} be the true value of the signal y at time $t+k$, and $\hat{y}_{t+k|t}$ be the value of y_{t+k} predicted at time t . The forecasts are computed at 4 pm (nominations deadline) for the day-ahead. With a market period duration of fifteen minutes, K is equal to 128. The PV forecasts are needed for lead times from $k = 33$ (00:00 to 00:15 am) to $k = K = 128$ (11:45 to 12:00 pm). Then,

$\hat{y}_{t+k|t}$ and y_{t+k} are assumed to be related by

$$\hat{y}_{t+k|t} = y_{t+k}(1 + \epsilon_k). \quad (10.18)$$

The error term ϵ_k is generated by the moving-average model [28, Chapter 3]

$$\epsilon_1 = \eta_1 \quad (10.19a)$$

$$\epsilon_k = \eta_k + \sum_{i=1}^{k-1} \alpha_i \eta_{k-i} \quad \forall k \in \{2, \dots, K\} \quad (10.19b)$$

with $\{\alpha_i\}_{i=1}^{K-1}$ scalar coefficients, $\{\eta_k\}_{k=1}^K$ independent and identically distributed sequences of random variables from a normal distribution $\mathcal{N}(0, \sigma)$. Thus, the variance of the error term is

$$\text{Var}(\epsilon_1) = \sigma^2 \quad (10.20a)$$

$$\text{Var}(\epsilon_k) = \left(1 + \sum_{i=1}^{k-1} \alpha_i^2\right) \sigma^2 \quad \forall k \in \{2, \dots, K\}. \quad (10.20b)$$

It is possible to simulate with this model an increase of the prediction error variance with the lead time k by choosing

$$\alpha_i = p^i \quad \forall i \in \{1, \dots, K-1\}. \quad (10.21)$$

(10.20a) becomes, $\forall k \in \{1, \dots, K\}$

$$\text{Var}(\epsilon_k) = \sigma^2 A_{\epsilon_k} \quad (10.22)$$

with A_{ϵ_k} defined $\forall k \in \{1, \dots, K\}$ by

$$A_{\epsilon_k} = \sum_{i=0}^{k-1} (p^2)^i = \frac{1 - (p^2)^k}{1 - p^2}. \quad (10.23)$$

Then, with $0 \leq p < 1$, it is possible to make the prediction error variance independent of the lead time as it increases. Indeed:

$$\lim_{k \rightarrow \infty} A_{\epsilon_k} = A_{\infty} = \frac{1}{1 - p^2}. \quad (10.24)$$

For instance, with $p = 0.9$ and $K = 128$, for $k \geq 33$, $A_{\epsilon_k} \approx A_\infty$ that is approximately 5.26. Thus, $\forall k \in \{33, \dots, K\}$

$$\text{Var}(\epsilon_k) \approx 5.26\sigma^2. \quad (10.25)$$

Finally, the σ value to set a maximum ϵ_{\max} with a high probability of 0.997, corresponding to a three standard deviation confidence interval from a normal distribution, is found by imposing $\epsilon_{\max} = 3\sqrt{\text{Var}(\epsilon_K)}$:

$$\sigma \approx \frac{\epsilon_{\max}}{3\sqrt{A_\infty}}, \quad (10.26)$$

with $\epsilon_{\max} = 0.25, 0.50, 0.75, 1.0$, $\sigma = 3.5, 7, 10.5, 14\%$.

Chapter 11

Capacity firming sizing

Overview

This Chapter proposes an approach to size a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. The main novelties in the CRE capacity framework context are three-fold.

1. First, a MIQP formulation is proposed to address the planning stage of the two-phase engagement control that is compatible with a scenario approach to approximate the Mixed-Integer Non-Linear Programming problem generated by the CRE non-convex penalty function. It is compared to the deterministic formulation using perfect knowledge of the future and PV point forecasts on empirical data from the PV production monitored on-site at the Liège University (ULiège), Belgium.
2. Second, a transparent and easily reproducible Gaussian copula methodology is implemented to generate PV scenarios based on the parametric PVUSA model using a regional climate model.
3. Finally, the sizing of the system is addressed by a grid search to approximate the optimal sizing for a given selling price using both the deterministic and stochastic approaches.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Bertrand Cornélusse, Xavier Fettweis, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Probabilistic forecasting for sizing in the capacity firming framework. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494947. URL <https://arxiv.org/abs/2106.02323>.

“Each player must accept the cards life deals him or her: but once they are in hand, he or she alone must decide how to play the cards in order to win the game.”

— Voltaire

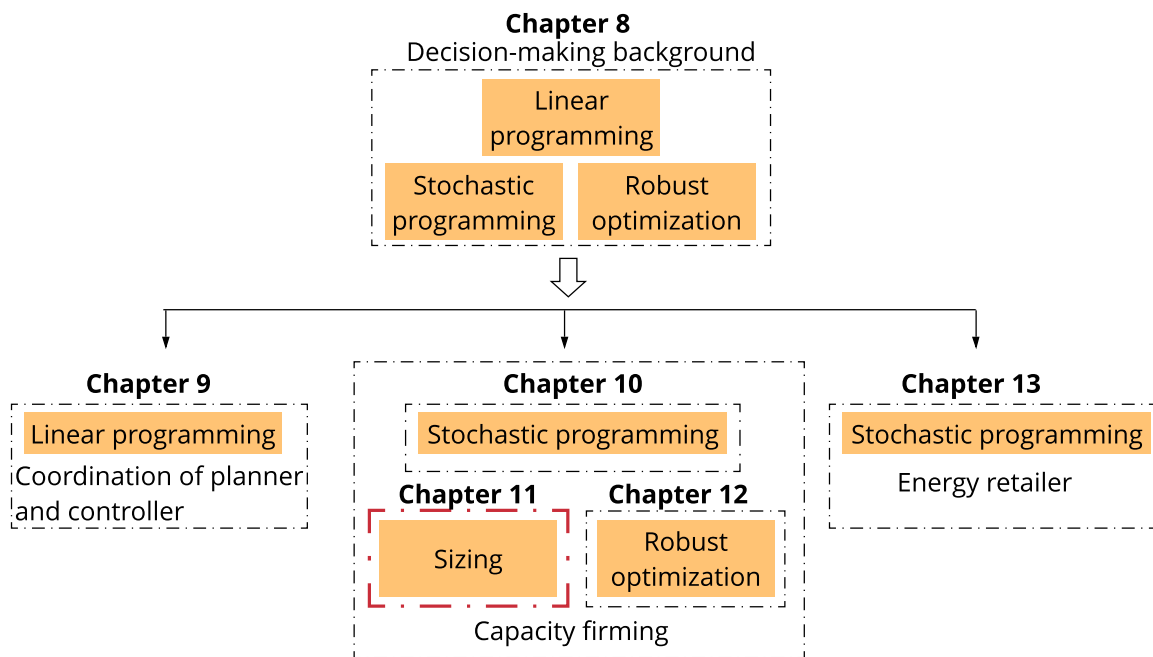


Fig. 11.1 Chapter 11 position in Part II.

The sizing has only been investigated by Haessig [80] in a similar context but with a wind farm. The main thread motivating the contribution of this Chapter is to extend Chapter 10 to address the sizing of a grid-connected PV plant and a BESS subject to grid constraints in this context. However, the sizing problem is difficult due to this two-phase engagement control with a day-ahead nomination and an intraday control to minimize deviations from the planning. A single sizing optimization problem would result in a non-linear bilevel optimization problem with an upper level, the sizing part, and a lower level, the two-phase engagement control. Thus, a grid search is conducted to approximate the optimal sizing for a given selling price using both the deterministic and stochastic approaches. In the two-phase engagement control, the PV uncertainty is considered at the planning stage by considering a stochastic approach that uses PV scenarios generated by a Gaussian copula methodology. The planner determines the engagement profile on a day-ahead basis given a selling price profile, the PV scenarios, and the system's current state, including the battery state of charge. Then, the engagement plan is sent to the grid operator and to the controller that computes every 15 minutes the production, *i.e.*, injection or withdrawal, and set-points, *e.g.*, BESS charge or discharge, until the end of the day. The optimization problems are formulated as Mixed-Integer Quadratic Problems (MIQP) with linear constraints.

The remainder of this Chapter is organized as follows. Section 11.1 defines the problem statement. Section 11.2 provides the PVUSA parametric point forecast model, and the Gaussian Copula approach to generate PV scenarios. Section 11.3 investigates the system sizing using both the deterministic and stochastic MIQP approaches. Finally, Section 11.4 summarizes the main findings and highlights ideas for further work. Note, the capacity firming process is described in Section 10.2 of Chapter 10.

11.1 Problem statement

In this Chapter, the problem formulation is almost strictly identical to Chapter 10. The only difference lies in the definition (10.2) of the penalty c . The penalty defined in the CRE specifications of the tender AO-CRE-ZNI 2019 published on *July* 12, 2019 is adopted to conduct the sizing. The optimization variables and the parameters are defined in Section 10.1.

11.1.1 Stochastic approach

A stochastic planner with a MIQP formulation and linear constraints is implemented using a scenario-based approach. The planner computes on a day-ahead basis the engagement plan x_t , $\forall t \in \mathcal{T}$ to be sent to the grid. The problem formulation is given by (10.15), where only short deviations are considered with the penalty defined in the CRE specifications of the tender AO-CRE-ZNI 2019 published on *July* 12, 2019. In compact form, the optimization problem is

$$\min \sum_{\omega \in \Omega} \alpha_{\omega} \sum_{t \in \mathcal{T}} \left[-\Delta t \pi_t y_{t,\omega} + \frac{\Delta t \pi_t}{P_c} d_{t,\omega}^- (d_{t,\omega}^- + 4pP_c) \right] \quad \text{s.t.:} \quad (11.1a)$$

$$x_t - x_{t-1} \leq \Delta X_t, \forall t \in \mathcal{T} \setminus \{1\} \quad (11.1b)$$

$$x_{t-1} - x_t \leq \Delta X_t, \forall t \in \mathcal{T} \setminus \{1\} \quad (11.1c)$$

$$x_t \leq x_t^{\max}, \forall t \in \mathcal{T} \quad (11.1d)$$

$$-x_t \leq -x_t^{\min}, \forall t \in \mathcal{T} \quad (11.1e)$$

$$-d_{t,\omega}^- \leq y_{t,\omega} - (x_t - pP_c), \forall t \in \mathcal{T} \quad (11.1f)$$

$$y_{t,\omega} \leq (x_t + pP_c), \forall t \in \mathcal{T} \quad (11.1g)$$

$$y_{t,\omega}^{\text{pv}} \leq \hat{y}_{t,\omega}^{\text{pv}}, \forall t \in \mathcal{T} \quad (11.1h)$$

$$y_{t,\omega}^{\text{cha}} \leq y_{t,\omega}^b y_{\max}^{\text{cha}}, \forall t \in \mathcal{T} \quad (11.1i)$$

$$y_{t,\omega}^{\text{dis}} \leq (1 - y_{t,\omega}^b) y_{\max}^{\text{dis}}, \forall t \in \mathcal{T} \quad (11.1j)$$

$$-s_{t,\omega} \leq -s^{\min}, \forall t \in \mathcal{T} \quad (11.1k)$$

$$s_{t,\omega} \leq s^{\max}, \forall t \in \mathcal{T} \quad (11.1l)$$

$$y_{t,\omega} = y_{t,\omega}^{\text{pv}} + y_{t,\omega}^{\text{dis}} - y_{t,\omega}^{\text{cha}}, \forall t \in \mathcal{T} \quad (11.1m)$$

$$y_{t,\omega} \leq y_t^{\max}, \forall t \in \mathcal{T} \quad (11.1n)$$

$$-y_{t,\omega} \leq -y_t^{\min}, \forall t \in \mathcal{T} \quad (11.1o)$$

$$s_{1,\omega} - s^{\text{ini}} = \Delta t \left(\eta^{\text{cha}} y_{1,\omega}^{\text{cha}} - \frac{y_{1,\omega}^{\text{dis}}}{\eta^{\text{dis}}} \right) \quad (11.1p)$$

$$s_{t,\omega} - s_{t-1,\omega} = \Delta t \left(\eta^{\text{cha}} y_{t,\omega}^{\text{cha}} - \frac{y_{t,\omega}^{\text{dis}}}{\eta^{\text{dis}}} \right), \forall t \in \mathcal{T} \setminus \{1\} \quad (11.1q)$$

$$s_{T,\omega} = s^{\text{end}} = s^{\text{ini}}. \quad (11.1r)$$

The optimization variables are x_t (engagement at the coupling point), $y_{t,\omega}$ (net power at the coupling point), $\forall \omega \in \Omega$, $d_{t,\omega}^-$ (underproduction), $y_{t,\omega}^{\text{pv}}$ (PV generation), $y_{t,\omega}^b$ (BESS binary variable), $y_{t,\omega}^{\text{cha}}$ (BESS charging power), $y_{t,\omega}^{\text{dis}}$ (BESS discharging power), and $s_{t,\omega}$ (BESS state of charge) (cf. the notation Section 10.1). The engagement

constraints are (11.1b)-(11.1e), where the ramping constraint on x_1 is deactivated to decouple consecutive days of simulation. The CRE non-convex piecewise quadratic penalty function is modeled by the constraints (11.1f) $\forall \omega \in \Omega$, that defines the variables $d_{t,\omega}^- \in \mathbb{R}_+$ to model the quadratic penalty for underproduction, and (11.1g) $\forall \omega \in \Omega$ forbidding overproduction that is non-optimal as curtailment is allowed [128]. The set of constraints that bound $y_{t,\omega}^{\text{pv}}$, $y_{t,\omega}^{\text{cha}}$, $y_{t,\omega}^{\text{dis}}$, and $s_{t,\omega}$ variables are (11.1h)-(11.1l) $\forall \omega \in \Omega$ where $\hat{y}_{t,\omega}^{\text{pv}}$ are PV scenarios, and $y_{t,\omega}^b$ are binary variables that prevent the BESS from charging and discharging simultaneously. The power balance equation and the production constraints are (11.1m) and (11.1n)-(11.1o) $\forall \omega \in \Omega$. The dynamics of the BESS state of charge is provided by constraints (11.1p)-(11.1r) $\forall \omega \in \Omega$. Note, the parameters s^{end} and s^{ini} are introduced to decouple consecutive days of simulation.

11.1.2 Deterministic approach

The deterministic (D) formulation of the planner is a specific case of the stochastic formulation by considering only one scenario where $y_{t,\omega}^{\text{pv}}$ become \hat{y}_t^{pv} , PV point forecasts. The deterministic formulation with perfect forecasts (D^*) is D with $\hat{y}_t^{\text{pv}} = y_t^{\text{pv,m}} \forall t \in \mathcal{T}$. For both the deterministic planners D and D^* , the optimization variables are x_t , y_t , d_t^- , y_t^{pv} , b_t , y_t^{cha} , y_t^{dis} , and s_t .

11.1.3 Oracle controller

The oracle controller is an ideal real-time controller that assumes perfect knowledge of the future by using as inputs the engagement profile to compute the set-points, maximize the revenues and minimize the deviations from the engagements. The oracle controller is D^* where the engagements are parameters.

11.2 Forecasting methodology

The Gaussian copula approach has been widely used to generate wind and PV scenarios in power systems [145, 72]. However, to the best of our knowledge, there is almost no guidance available on which copula family can describe correlated variations in PV generation [72], hence the Gaussian copula family is selected instead of copulas like Archimedean or Elliptical.

11.2.1 Gaussian copula-based PV scenarios

In this section, let t be the current time index, k be the lead time of the prediction, $Z = \{Z_1, \dots, Z_T\}$ a multivariate random variable, $F_{Z_k}(\cdot)$, $k = 1, \dots, T$ the marginal cumulative distribution functions, and R_Z the correlation matrix. The goal is to generate samples of Z . The Gaussian copula methodology consists of generating a trajectory $u = [u_1, \dots, u_T]$ from the multivariate Normal distribution $\mathcal{N}(0, R_Z)$. Then, to transform each entry u_k through the standard normal cumulative distribution function $\phi(\cdot)$: $\tilde{u}_k = \phi(u_k)$, $k = 1, \dots, T$, and finally, to apply to each entry \tilde{u}_k the inverse marginal cumulative distribution function of Z_k : $z_k = F_{Z_k}^{-1}(\tilde{u}_k)$, $k = 1, \dots, T$. In our case, Z is defined as the error between the PV measurement x and the PV point forecast \hat{y}^{PV}

$$Z_k = y_{t+k}^{\text{PV,m}} - \hat{y}_{t+k|t}^{\text{PV}} \quad k = 1, \dots, T. \quad (11.2)$$

$\hat{F}_{Z_k}(\cdot)$ and \hat{R}_Z are estimated from the data, and following the methodology described above, a PV scenario i at time t is generated for time $t + k$

$$\hat{y}_{t+k|t}^{\text{PV},(i)} = \hat{y}_{t+k|t}^{\text{PV}} + z_k^i \quad k = 1, \dots, T. \quad (11.3)$$

The PV point forecasts $\hat{y}_{t+k|t}^{\text{PV}}$ are computed by using the PVUSA model presented in the following Section.

11.2.2 PV point forecast parametric model

A PV plant can be modeled using the well-known PVUSA parametric model [46], which expresses the instantaneous generated power as a function of irradiance and air temperature

$$y_t^{\text{PV,m}} = aI_t + bI_t^2 + cI_tT_t, \quad (11.4)$$

where $y_t^{\text{PV,m}}$, I_t and T_t are the generated power, irradiance and air temperature at time t , respectively, and $a > 0$, $b < 0$, $c < 0$ are the PVUSA model parameters. These parameters are estimated following the algorithm of Bianchini et al. [19] that efficiently exploits only the power generation measurements and the theoretical clear-sky irradiance and is characterized by lightweight computational effort. The same implementation of the algorithm is used with a sliding window of 12 hours. The parameters reached the steady-state values in 50 days on the Uliège case study, described in the following

Section, with

$$[\tilde{a}, \tilde{b}, \tilde{c}]^\top = [0.573, -7.68 \cdot 10^{-5}, -1.86 \cdot 10^{-3}]^\top. \quad (11.5)$$

The weather hindcasts from the MAR regional climate model [59], provided by the Laboratory of Climatology of the Liège University, are used as inputs of the parametric PVUSA model. The ERA5¹ reanalysis database forces the MAR regional climate to produce weather hindcasts. Finally, the PV point forecasts are computed

$$\hat{g}_{t+k|t}^{\text{PV}} = \tilde{a}\hat{I}_{t+k|t} + \tilde{b}\hat{I}_{t+k|t}^2 + \tilde{c}\hat{I}_{t+k|t}\hat{T}_{t+k|t}, \quad \forall k = k_1, \dots, k_T, \quad (11.6)$$

and use as inputs to generate PV scenarios following the Gaussian copula approach.

11.2.3 PV scenarios

The Uliège case study is composed of a PV generation plant with an installed capacity of 466.4 kWp. The simulation dataset is composed of August 2019 to December 2019, 151 days in total, with a total production of 141.3 MWh. The Uliège PV generation is monitored on a minute basis and is resampled to 15 minutes. The set of PV scenarios is generated using the Gaussian copula approach based on the PVUSA point forecasts. Figure 11.2 illustrates a set of 5 PV scenarios on four days of the dataset. The PV scenarios are used as input of a stochastic optimization problem.

11.3 Sizing study

The goal is to determine the optimal sizing of the BESS and PV for a given selling price and its related net revenue over the lifetime project to bid at the tendering stage optimally.

11.3.1 Problem statement and assumptions

The ratio $r_{\bar{S}} = \frac{\bar{S}}{P_c}$, BESS maximum capacity over the PV installed capacity, is introduced to model several BESS and PV system configurations. The total exports, imports, deviation costs, number of charging and discharging cycles for several $r_{\bar{S}}$ and selling prices are computed. Based on the BESS and PV CapEx (Capital Expenditure)

¹The fifth-generation European Centre for Medium-Range Weather Forecasts atmospheric reanalysis of the global climate.

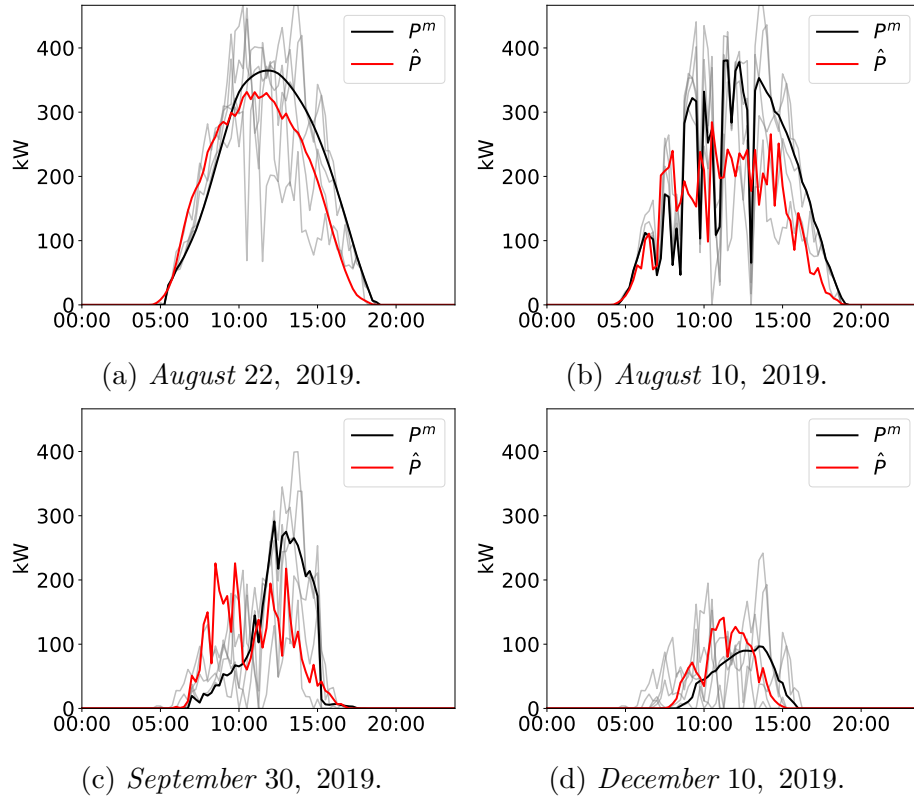


Fig. 11.2 5 PV scenarios (gray), PVUSA point forecasts (red), and PV measurements (black).

I and OpEx (Operational Expenditure) O&M, the LCOE (Levelized Cost Of Energy) is calculated

$$\text{LCOE} = \frac{\text{CRF} \cdot I + \text{O\&M} + W + C}{E}, \quad (11.7)$$

with CRF the Capital Recovery Factor (or Annuity Factor), and E, W, C, the annual export, withdrawal, and deviation costs, respectively. Then, the net revenue over the lifetime project is defined as the annual gross revenue R divided by the annual export to the grid minus the LCOE

$$\text{net}(\pi, r_{\bar{S}}) := \frac{R}{E} - \text{LCOE}. \quad (11.8)$$

The higher the net revenue, the more profitable the system. Section 11.3.2 details the LCOE definition and the assumptions to establish (11.7) and (11.8). Finally, the

optimal sizing for a given selling price is provided by

$$r_{\bar{S}}^*(\pi) = \arg \max_{r_{\bar{S}} \in \mathcal{R}_{\bar{S}}} \text{net}(\pi, r_{\bar{S}}), \quad (11.9)$$

with $\mathcal{R}_{\bar{S}}$ the sizing space, and the sizing approach is depicted in Figure (11.3).

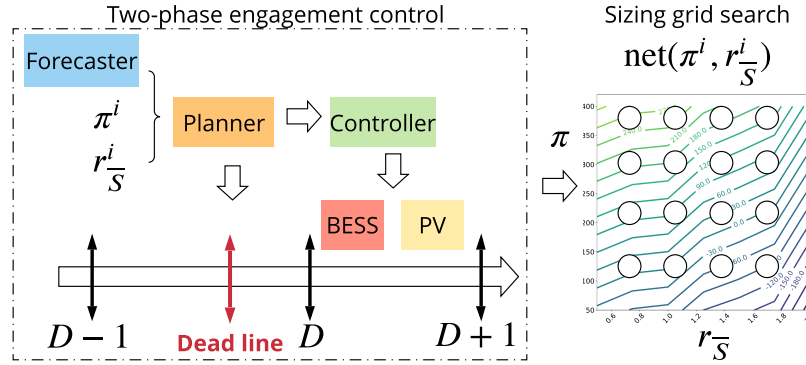


Fig. 11.3 Sizing approach.

11.3.2 Levelized cost of energy (LCOE)

The LCOE (€/ MWh) is "the unique break-even cost price where discounted revenues, price times quantities, are equal to the discounted net expenses"². It means the LCOE is the lower bound of the selling price to be profitable over the lifetime of the project and is defined as

$$\sum_{t=0}^n \frac{E_t \cdot \text{LCOE}}{(1+i)^t} = \sum_{t=0}^n \frac{\text{Expenses}_t}{(1+i)^t}, \quad (11.10)$$

with i (%) the discount rate, n (years) the lifetime of the project, and E_t (MWh) the annual energy at year t . The lifetime expenses comprise investment costs I that are the upfront investment (or CapEx) (€/kW for PV and €/kWh for BESS), operation and maintenance cost (or OpEx) O&M (€/kW for PV and €/kWh for BESS), the annual cost of the energy withdrawn from the grid W (€), and the annual deviation cost penalty C (€)

$$\text{LCOE} = \frac{\sum_{t=0}^n \frac{I_t + \text{O\&M}_t + W_t + C_t}{(1+i)^t}}{\sum_{t=0}^n \frac{E_t}{(1+i)^t}}. \quad (11.11)$$

²Annex II Metrics & Methodology of the AR5 IPCC report.

We assume the annual cost of the energy E_t (W_t) exported (withdrawn), the annual OpEx $O\&M_t$, and the annual deviation cost C_t are constant during the lifetime of the project: E, W, C and $O\&M \forall t > 0$. The investment costs I are the sum of all capital expenditures needed to make the investment fully operational discounted to $t = 0$. Thus, $I_t = 0 \forall t > 0$ and $I_0 = I$. Finally, the system is operational at year $t = 1$: $E_0 = 0$, $O\&M_0 = 0$, $W_0 = 0$, and $C_0 = 0$. (11.11) becomes

$$\text{LCOE} = \frac{I + (O\&M + W + C) \sum_{t=1}^n \frac{1}{(1+i)^t}}{E \sum_{t=1}^n \frac{1}{(1+i)^t}}, \quad (11.12)$$

that is re-written

$$\text{LCOE} = \frac{\text{CRF} \cdot I + O\&M + W + C}{E}, \quad (11.13)$$

with $\text{CRF} = (\sum_{t=1}^n \frac{1}{(1+i)^t})^{-1} = \frac{i}{1-(1+i)^{-n}}$. The gross revenue is the energy exported to the grid that is remunerated at the selling price. The annual gross revenue R is assumed to be constant over the project lifetime. Then, the LCOE can be compared to R divided by the annual export E to assess the financial viability of the project by calculating (11.8).

11.3.3 Case study description

The ULiège case study comprises a PV generation plant with an installed capacity P_c of 466.4 kW. The period from August 2019 to December 2019, 151 days in total, composes the dataset. The PV generation is monitored on a minute basis and resampled to 15 minutes.

The simulation parameters of the planners and the oracle controller are identical. The planning and controlling periods are $\Delta t = 15$ minutes. The peak hours are between 7 pm and 9 pm (UTC+0). The specifications of the tender AO-CRE-ZNI 2019 published on *July 12, 2019* define the ramping power constraint on the engagements $\Delta X = 7.5\%P_c$ ($15\%P_c$) during off-peak (peak) hours. The lower bound on the engagement is $x_t^{\min} = -5\%P_c$ ($x_t^{\min} = 20\%P_c$) during off-peak (peak) hours. The lower bound on the production is $y_t^{\min} = -5\%P_c$ ($y_t^{\min} = 15\%P_c$) during off-peak (peak) hours. The upper bounds on the engagement and production are $x_t^{\max} = y_t^{\max} = P_c$. Finally, the engagement deadband is 5% of P_c .

The Python Pyomo³ 5.6.7 library is used to implement the algorithms in Python 3.7. IBM ILOG CPLEX Optimization Studio⁴ 12.9 is used to solve all the mixed-integer quadratic optimization problems. Numerical experiments are performed on an Intel Core i7-8700 3.20 GHz based computer with 12 threads and 32 GB of RAM running on Ubuntu⁵ 18.04 LTS. The average computation time per optimization problem of the S planner with $\#\Omega = 20$ scenarios is 3 (s) for an optimization problem with 15 000 variables and 22 000 constraints.

11.3.4 Sizing parameters

The BESS and PV CAPEX are 300 €/kWh and 700 €/kW, the BESS and PV OPEX are 1% of the CAPEX, the project lifetime is 20 years, and the weighted average cost of capital is 5%. The BESS lifetime in terms of complete charging and discharging cycles is 3 000. The BESS is assumed to be capable of fully charging or discharging in one hour $y_{\max}^{\text{cha}} = y_{\max}^{\text{dis}} = s^{\max}/1$, with charging and discharging efficiencies $\eta^{\text{cha}} = \eta^{\text{dis}} = 0.95$. Each simulation day is independent with a discharged battery at the first and last period to its minimum capacity $s^{\text{ini}} = s^{\text{end}} = 10\%s^{\max}$. The BESS minimum (s^{\min}) and maximum (s^{\max}) capacities are 10% and 90% of the total BESS storage capacity \bar{S} . The sizing space is a grid composed of 56 values with $\mathcal{R}_{\bar{S}} = \{0.5, 0.75, 1, 1.25, 1.5, 1.75, 2\}$ and $\mathcal{P} = \{50, 100, 150, 200, 250, 300, 350, 400\}$.

11.3.5 Sizing results

The D^* , D , and $S^{\Omega=20}$ planners are used with the oracle controller over the simulation dataset. E, W, C, and the number of complete charging and discharging cycles are calculated by extrapolating the 151 days to one year. Figure 11.4 provides the grid search sizing results for the three planners. For a given selling price, the net is maximal when $r_{\bar{S}} = 0.5$. When $r_{\bar{S}}$ increases, the BESS is more and more used to withdraw and export during peak hours. It leads to an increase in the number of charging/discharging cycles that implies an increase of the number of BESS required during the project lifetime and consequently an increase of the BESS CAPEX. As the BESS CAPEX mainly drives the LCOE, $\frac{R}{E}$ is not capable of compensating the LCOE increase resulting in a net decrease. The number of charging and discharging cycles is approximately the

³[www.http://www.pyomo.org/](http://www.pyomo.org/)

⁴<https://www.ibm.com/products/ilog-cplex-optimization-studio>

⁵<https://ubuntu.com>

same for both the D and S planners, independently of the selling price, and rises from 4 700 with $r_{\bar{S}} = 0.5$ to 13 000 with $r_{\bar{S}} = 2$.

The differences between D^* , D , and S planners are minor. D^* tends to overestimate the net by underestimating the LCOE (underestimating the deviation, BESS CAPEX, and withdrawal costs) and overestimating $\frac{R}{E}$. However, the minimal selling price to be profitable with $r_{\bar{S}} = 0.5$, is approximately 80 €/ MWh for all planners as shown by Figure 11.4. Then, the higher the selling price, the higher the net. In the CRE specifications, the best tender is mainly selected based on the selling price criterion. A trade-off should be reached between the net and the selling price to be selected.

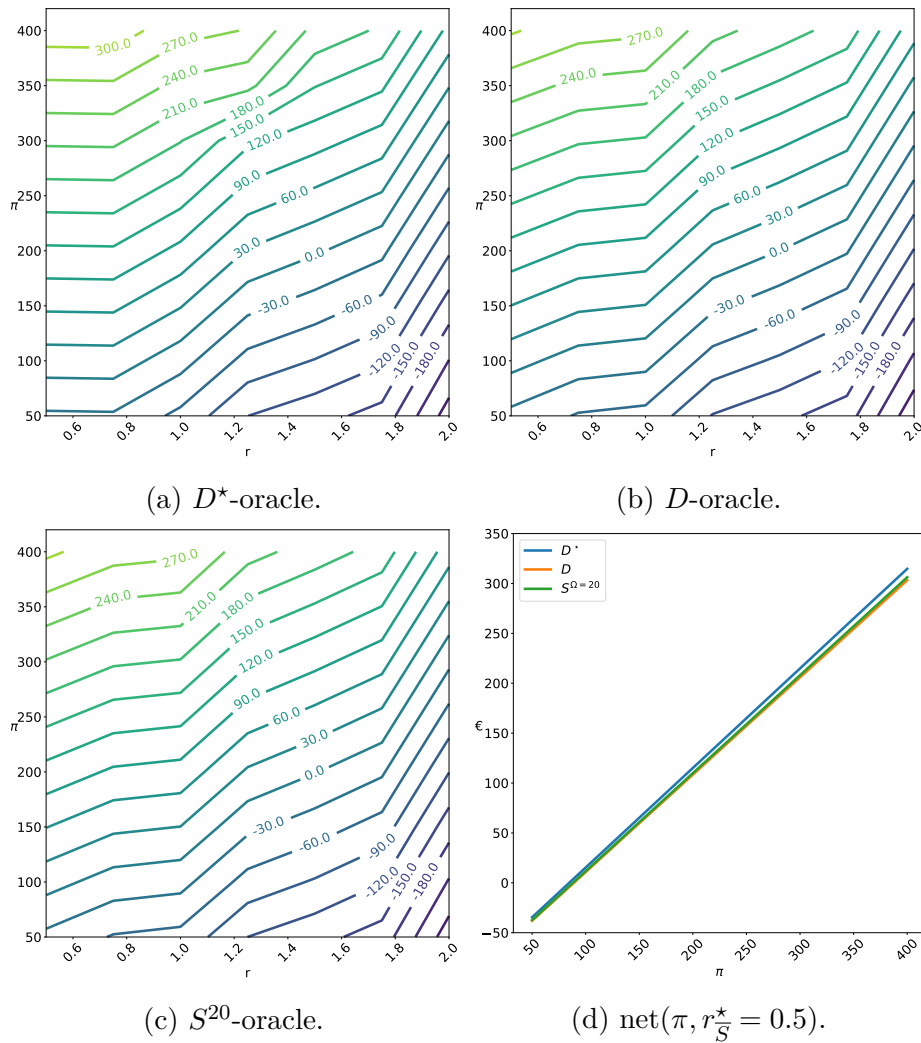


Fig. 11.4 Sizing results: $\text{net}(\pi, r_{\bar{S}}^*)$ (€/ MWh).

This sizing study seems to indicate that it is not very sensitive to the control policy, *i.e.*, deterministic with perfect knowledge, deterministic with point forecasts,

and stochastic with scenarios. However, it may be dangerous not considering the uncertainty at the sizing stage and could lead to overestimating the system performance and underestimating the sizing. Indeed, the two-phase engagement control net revenues are similar between planners explaining the minor differences in terms of sizing. Two main limitations could explain this result. First, large deviations (15-20%) from the engagement plan at the control step occur rarely. Indeed, the oracle controller may compensate for inadequate planning and limits the deviations. A more realistic controller with point forecasts should be considered. Second, when such deviations occur, they are usually within the tolerance where there is no penalty. Furthermore, when the deviations are outside, the penalty is relatively small in comparison with the gross revenue. A sensitivity analysis of the numerical settings of the CRE specifications should be performed.

11.4 Conclusions and perspectives

The key idea of this Chapter is to propose a methodology to size the PV and BESS in the context of the capacity firming framework. Indeed, the two-phase engagement control cannot easily be modeled as a single sizing optimization problem. Such an approach would result in a non-linear bilevel optimization problem challenging to solve with an upper level, the sizing part, and a lower level, the two-phase engagement control. The two-phase engagement control is decomposed into two steps: computing the day-ahead engagements, then recomputing the set-points in real-time to minimize the deviations from the engagements. The CRE non-convex penalty function is modeled by a threshold-quadratic penalty that is compatible with a scenario approach. The stochastic formulation using a scenario approach is compared to the deterministic formulation. The PV scenarios are generated using a Gaussian copula methodology and PV point forecasts computed with the PVUSA model. The minimal selling price to be profitable, on this dataset, in the context of the capacity firming framework is approximately 80 €/ MWh with a BESS having a maximal capacity, fully charged or discharged in one hour, of half the total PV installed power. The sizing study indicates that it is not very sensitive to the control policy. The differences are minor between the deterministic with perfect knowledge (or point forecasts) and stochastic with scenarios strategies. However, further investigations are required to implement a more realistic controller that uses intraday point forecasts and conduct a sensitivity analysis on the simulation parameters.

Several extensions are under investigation.

-
- A PV generation methodology that is less dependent on the PV point forecasts and considers the PV power's error dependency should be implemented.
 - PV scenarios clustering and reduction techniques could be considered to select relevant PV scenarios and improve the stochastic planner results.
 - A sizing formulation as a single optimization problem with the PV and BESS capacities as variables. It allows to compute the optimum directly and avoid doing a grid search. However, this formulation is not trivial due to the specific two-phase engagement control of the capacity firming framework.
 - Finally, the BESS aging process could be modeled in the sizing study. A dataset with at least a full year of data should be considered to consider the PV seasonality fully.

Chapter 12

Capacity firming using a robust approach

Overview

This Chapter proposes a robust approach to address the energy management of a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. It is an extension of Chapter 10 by considering another optimization formulation to handle the PV uncertainty. The main contributions are two-fold:

1. The core contribution is applying the robust optimization framework to the capacity firming market in a tractable manner thanks to a Benders decomposition of the optimization problem and a warm start of the algorithm. In addition, a dynamic risk-averse parameters selection taking advantage of the quantile forecast distribution is proposed.
2. The secondary contribution is the use of the Normalizing Flows, which is a new advanced forecasting technique, to provide the uncertainty estimation in the form of PV quantiles for the robust planner. To the best of our knowledge, it is the first study to use NFs in a power system application.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Colin Cointe, Antoine Wehenkel, Antonio Sutera, Xavier Fetweis, and Bertrand Cornélusse. A probabilistic forecast-driven strategy for a risk-aware participation in the capacity firming market. *IEEE Transactions on Sustainable Energy*, pages 1–1, 2021. doi: 10.1109/TSTE.2021.3117594. URL <https://arxiv.org/abs/2105.13801>.

“ You have to start with the truth. The truth is the only way that we can get anywhere. Because any decision-making that is based upon lies or ignorance can't lead to a good conclusion.”

— Julian Assange

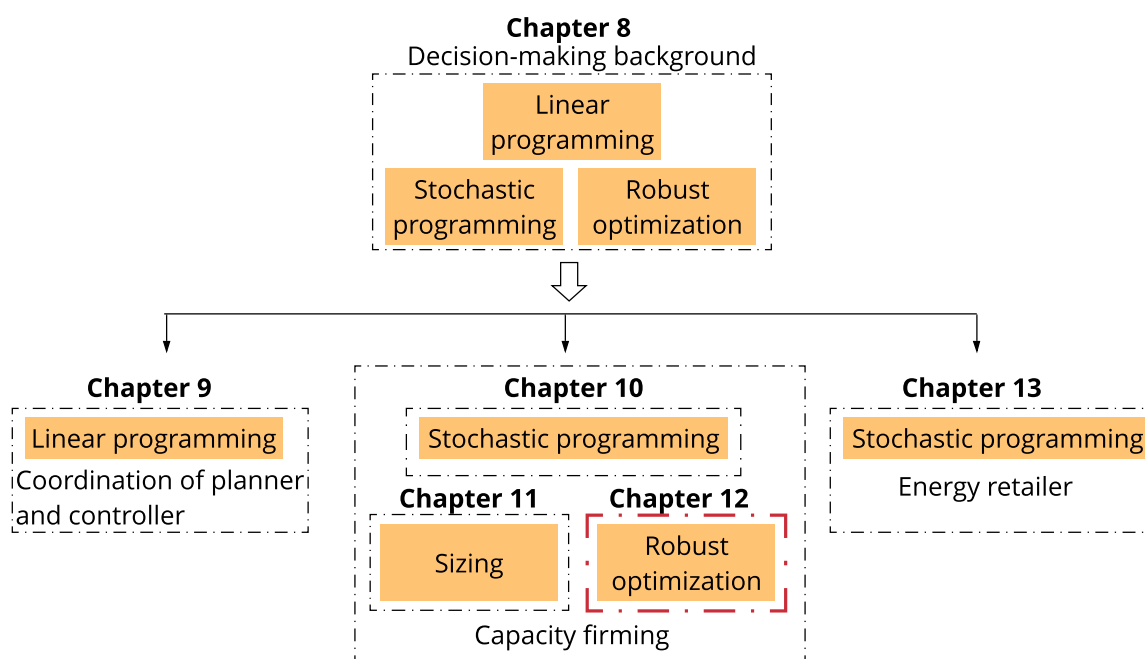


Fig. 12.1 Chapter 12 position in Part II.

There are several approaches to deal with renewable energy uncertainty. One way is to consider a two-stage stochastic programming approach [22]. It has already been applied to the capacity firming framework [53, 127, 80, 138]. The generation uncertainty is captured by a set of scenarios modeling possible realizations of the power output. However, this approach has three drawbacks. First, the problem size and computational requirement increase with the number of scenarios, and a large number of scenarios are often required to ensure the good quality of the solution. Second, the accuracy of the algorithm is sensitive to the scenario generation technique. Finally, it may be challenging to identify an accurate probability distribution of the uncertainty. Another option is to consider robust optimization (RO) [11, 17], applied to unit commitment by [18, 95], and in the capacity firming setting [127]. RO accounts for the worst generation trajectory to hedge the power output uncertainty, where the uncertainty model is deterministic and set-based. Indeed, the RO approach puts the random problem parameters in a predetermined uncertainty set containing the worst-case scenario. It has two main advantages [18]: (1) it only requires moderate information about the underlying uncertainty, such as the mean and the range of the uncertain data; (2) it constructs an optimal solution that immunizes against all realizations of the uncertain data within a deterministic uncertainty set. Therefore, RO is consistent with the risk-averse fashion way to operate power systems. However, the RO version of a tractable optimization problem may not itself be tractable, and some care must be taken in choosing the uncertainty set to ensure that tractability is preserved.

Traditionally, a two-stage RO model is implemented for the unit commitment problem in the presence of uncertainty. However, it is challenging to compute and often NP-hard. Two classes of cutting plane strategies have been developed to overcome the computational burden. The Benders-dual cutting plane (BD) algorithms are the most used and seek to derive exact solutions in the line of Benders' decomposition [12] method. They decompose the overall problem into a master problem involving the first-stage commitment decisions at the outer level and a sub-problem associated with the second-stage dispatch actions at the inner level. Then, they gradually construct the value function of the first-stage decisions using dual solutions of the second-stage decision problems [18, 95]. In contrast, the column-and-constraint generation (CCG) procedure, introduced by [189, 183] does not create constraints using dual solutions of the second-stage decision problem. Instead, it dynamically generates constraints with recourse decision variables in the primal space for an identified scenario. The generated variables and constraints in the CCG procedure are similar to those in the

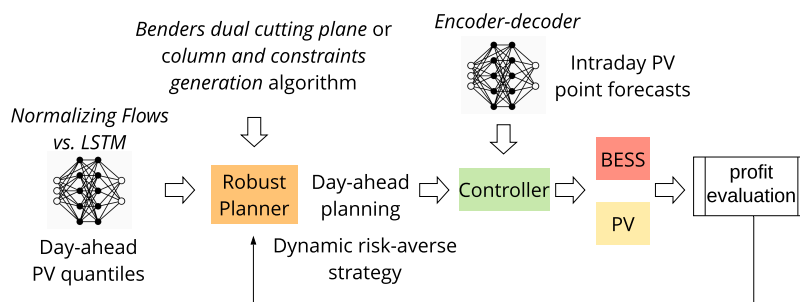


Fig. 12.2 Forecast-driven robust optimization strategy.

deterministic equivalent of a two-stage stochastic programming model. The BD and CCG algorithms have not been compared in the capacity firming framework to the best of our knowledge.

This Chapter proposes a reliable and computationally tractable probabilistic forecast-driven robust optimization strategy. It can use either a BD or CCG algorithm in the capacity firming framework, depicted in Figure 12.2. Our work goes several steps further than [127]. The main contributions of this Chapter are three-fold:

1. The core contribution is applying the robust optimization framework to the capacity firming market in a tractable manner by using a Benders decomposition. The non-linear robust optimization problem is solved both using the Benders-dual cutting plane and the column-and-constraint generation algorithms. To the best of our knowledge, it is the first time that a comparison of these algorithms is performed in the capacity firming framework. In addition, the convergence of the BD algorithm is improved with a warm-start procedure. It consists of building an initial set of cuts based on renewable generation trajectories assumed to be close to the worst-case scenario. The results of both the CCG and BD two-stage RO planners are compared to the deterministic planner using perfect knowledge of the future, the nominal point forecasts, *i.e.*, the baseline to outperform, and the quantiles (a conservative approach). The case study is the photovoltaic (PV) generation monitored on-site at the University of Liège (ULiège), Belgium.
2. Second, a dynamic risk-averse parameters selection taking advantage of the quantile forecast distribution is investigated and compared to a strategy with fixed risk-averse parameters.
3. Finally, the *normalizing flows* (NFs) is implemented. A new class of probabilistic generative models has gained increasing interest from the deep learning community in recent years. NFs are used to compute day-ahead quantiles of renewable

generation for the robust planner. Then, an encoder-decoder architecture forecasting model [27] computes the intraday point forecasts for the controller. To the best of our knowledge, it is the first study to use NFs in a power system application.

In addition to these contributions, this study also provides open access to the Python code¹ to help the community to reproduce the experiments. The rest of this Chapter is organized as follows. Section 12.1 provides the mathematical formulations of the robust and deterministic planners. Section 12.2 develops the Benders-dual cutting plane algorithm. The case study and computational results are presented in Section 12.3. Finally, Section 12.4 concludes this study and draws some perspectives of future works. Note, the capacity firming process is described in Section 10.2 of Chapter 10, and Section 4.5 in Chapter 4 introduces the forecasting techniques and proposes a quality evaluation.

12.1 Problem formulation

For the sake of simplicity in this Chapter, the penalty c defined in (10.2) is assumed to be symmetric, convex, and piecewise-linear

$$c(x_t, y_t^m) = \pi \Delta t \beta \left(\max \left(0, |x_t - y_t^m| - pP_c \right) \right), \quad (12.1)$$

with β a penalty factor. Note, in this Chapter, both the import and export are considered in contrast to Chapter 10. A two-stage robust optimization formulation is built to deal with the engagement for the uncertain renewable generation that is modeled with an uncertainty set. The deterministic and robust formulations of the planner are presented in Sections 12.1.1 and 12.1.2. The robust optimization problem with recourse has the general form of a min-max-min optimization problem. The uncertainty set is defined by quantiles forecasts and a budget of uncertainty Γ . Section 12.1.3 uses the dual of the inner problem to formulate a min-max optimization problem. Finally, Section 12.1.4 presents the formulation of the controller. The optimization variables and the parameters are defined in Section 10.1.

¹<https://github.com/jonathandumas/capacity-firming-ro>

12.1.1 Deterministic planner formulation

The deterministic formulation of the planner is provided in Section 10.3 of Chapter 10. The only difference lies in the definition of the penalty c that is assumed to be symmetric, convex, and piecewise-linear instead of a quadratic penalty in (10.5). Therefore, in this Chapter, the objective function J of the deterministic planner to minimize is

$$J(x_t, y_t) = \sum_{t \in \mathcal{T}} \pi_t \Delta t [-y_t + \beta(d_t^- + d_t^+)]. \quad (12.2)$$

The deterministic formulation is the following Mixed-Integer Linear Program (MILP)

$$\begin{aligned} \min_{x_t \in \mathcal{X}, y_t \in \Omega(x_t, \hat{y}_t^{\text{pv}})} J(x_t, y_t) \\ \mathcal{X} = \{x_t : (10.6)\} \\ \Omega(x_t, \hat{y}_t^{\text{pv}}) = \{y_t : (10.7) - (10.11)\} \end{aligned} \quad (12.3)$$

where \mathcal{X} and $\Omega(x_t, \hat{y}_t^{\text{pv}})$ are the sets of feasible engagements x_t and dispatch solutions y_t , respectively. The optimization variables of (12.3) are the engagement variables x_t , the dispatch variables y_t (the net power at the grid connection point), y_t^{dis} (BESS discharging power), y_t^{cha} (BESS charging power), s_t (BESS state of charge), y_t^b (BESS binary variables), y_t^{pv} (renewable generation), and d_t^-, d_t^+ (deviation variables) (cf. the notation Section 10.1).

12.1.2 Robust planner formulation

The uncertain renewable generation \hat{y}_t^{pv} of (10.11) is assumed to be within an interval $\mathcal{U} = [u_t^{\min}, u_t^{\max}]$ that can be obtained based on the historical data or an interval forecast composed of quantiles. In the following \hat{y}_t^{pv} is replaced by u_t in (10.11). The proposed two-stage robust formulation of the capacity firming problem consists of minimizing the objective function over the worst renewable generation trajectory

$$\max_{u_t \in \mathcal{U}} \left[\min_{x_t \in \mathcal{X}, y_t \in \Omega(x_t, u_t)} J(x_t, y_t) \right], \quad (12.4)$$

that is equivalent to

$$\min_{x_t \in \mathcal{X}} \left[\max_{u_t \in \mathcal{U}} \min_{y_t \in \Omega(x_t, u_t)} J(x_t, y_t) \right]. \quad (12.5)$$

The worst-case dispatch cost has a max-min form, where

$$\min_{y_t \in \Omega(x_t, u_t)} J(x_t, y_t)$$

determines the economic dispatch cost for a fixed engagement and a renewable generation trajectory, which is then maximized over the uncertainty set \mathcal{U} .

In the capacity firming framework, where curtailment is allowed, the uncertainty interval consists only in downward deviations $[u^{\min}, \hat{y}_t^{\text{pv},(0.5)}]$, with $\hat{y}_t^{\text{pv},(0.5)}$ the PV 50% quantile forecast.

Demonstration 1. *Let consider $\mathcal{U}_1 = [u_t^{\min}, \hat{y}_t^{\text{pv},(0.5)}]$, $\mathcal{U}_2 = [\hat{y}_t^{\text{pv},(0.5)}, u_t^{\max}]$, $u_t^1 \in \mathcal{U}_1$, $u_t^2 \in \mathcal{U}_2$, $y_t^1 \in \Omega_1 = \Omega(x_t, u_t^1 \in \mathcal{U}_1)$, and $y_t^2 \in \Omega_2 = \Omega(x_t, u_t^2 \in \mathcal{U}_2)$. Then, for a given engagement x_t*

$$\max_{u_t \in \mathcal{U}} \min_{y_t \in \Omega(x_t, u_t)} J(x_t, y_t) = \max \left[\max_{u_t^1 \in \mathcal{U}_1} \min_{y_t^1 \in \Omega_1} J(x_t, y_t^1), \max_{u_t^2 \in \mathcal{U}_2} \min_{y_t^2 \in \Omega_2} J(x_t, y_t^2) \right]. \quad (12.6)$$

The only difference between Ω_1 and Ω_2 is provided by (10.11) where $y_{t,1}^{\text{pv}} \leq u_t^1$ and $y_{t,2}^{\text{pv}} \leq u_t^2$. By definition of the uncertainty sets $u_t^1 \leq u_t^2 \forall t \in \mathcal{T}$, and it is straightforward that $y_{t,1}^{\text{pv}} \leq u_t^2$. In addition, the variables y_t^1 satisfy all the other constraints of $\Omega_2 \forall t \in \mathcal{T}$. Therefore, $y_t^1 \in \Omega_2 \forall t \in \mathcal{T}$, and $\Omega_1 \subseteq \Omega_2$. Thus, for a given engagement x_t , $\forall u_1 \in \mathcal{U}_1$, and $\forall u_2 \in \mathcal{U}_2$

$$\min_{y_t^2 \in \Omega_2} J(x_t, y_t^2) = J_2^* \leq \min_{y_t^1 \in \Omega_1} J(x_t, y_t^1) = J_1^*. \quad (12.7)$$

Finally, $\max[J_2^*, J_1^*] = J_1^*$. It means the worst-case is in Ω_1 that corresponds to \mathcal{U}_1 .

In addition, the worst generation trajectories, in robust unit commitment problems, are achieved when the uncertain renewable generation u_t reaches the lower or upper bounds of the uncertainty set [189, Proposition 2]. Thus, the uncertainty set at t is composed of two values and $u_t \in \{u^{\min}, \hat{y}_t^{\text{pv},(0.5)}\}$.

Following [18, 95], to adjust the degree of conservatism, a budget of uncertainty Γ taking integer values between 0 and 95 is employed to restrict the number of periods that allow u_t to be far away from its nominal value, *i.e.*, deviations are substantial. Therefore, the uncertainty set of renewable generation \mathcal{U} is defined as follows

$$\mathcal{U} = \left\{ u_t \in \mathbb{R}^T : \sum_{t \in \mathcal{T}} z_t \leq \Gamma, z_t \in \{0; 1\}, u_t = \hat{y}_t^{\text{pv},(0.5)} - z_t u_t^{\min} \forall t \in \mathcal{T} \right\}, \quad (12.8)$$

where $u_t^{\min} = \hat{y}_t^{\text{pv},(0.5)} - \hat{y}_t^{\text{pv},(q)}$, with $0 \leq q \leq 0.5$. When $\Gamma = 0$, the uncertainty set $\mathcal{U} = \{\hat{y}_t^{\text{pv},(0.5)}\}$ is a singleton, corresponding to the nominal deterministic case. As Γ increases the size of \mathcal{U} enlarges. This means that a larger total deviation from the expected renewable generation is considered, so that the resulting robust solutions are more conservative and the system is protected against a higher degree of uncertainty. When $\Gamma = T$, \mathcal{U} spans the entire hypercube defined by the intervals for each u_t .

12.1.3 Second-stage planner transformation

The robust formulation (12.5) consists of solving a min-max-min problem, which cannot be solved directly by a commercial software such as CPLEX or GUROBI. A scenario-based approach, *e.g.*, enumerating all possible outcomes of u_t that could lead to the worst-case scenario for the problem, results in at least 2^Γ possible trajectories². Thus, to deal with the huge size of the problem a Benders type decomposition algorithm is implemented.

Constraints (10.7a)-(10.7b) make the dispatch problem a MILP, for which a dual formulation cannot be derived. In view of this, following [95], the constraints (10.7a)-(10.7b) are relaxed (the convergence of the relaxed dispatch problem is discussed in Section 12.2.2). Then, by applying standard tools of duality theory in linear programming, the constraints and the objective function of the dual of the dispatch problem are derived. The dual of the feasible set $\Omega(x_t, u_t)$, with (10.7a)-(10.7b) relaxed, provides the dual variables ϕ_t and the following objective

$$G(x_t, u_t, \phi_t) = \sum_{t \in \mathcal{T}} \left[\phi_t^{\text{cha}} y_{\max}^{\text{cha}} + \phi_t^{\text{dis}} y_{\max}^{\text{dis}} - \phi_t^{s^{\min}} s^{\min} + \phi_t^{s^{\max}} s^{\max} - \phi_t^{y^{\min}} y^{\min} + \phi_t^{y^{\max}} y^{\max} + \phi^{s^{\text{ini}}} s^{\text{ini}} + \phi^{s^{\text{end}}} s^{\text{end}} - \phi_t^{d^-} (x_t - pP_c) + \phi_t^{d^+} (x_t + pP_c) + \phi_t^{y^{\text{pv}}} u_t \right]. \quad (12.9)$$

Then, the dual of the dispatch problem $\min_{y_t \in \Omega(x_t, u_t)} J(x_t, y_t)$ is

$$\max_{\phi_t \in \Phi} G(x_t, u_t, \phi_t) \quad (12.10)$$

$$\Phi = \left\{ (12.11a) - (12.11k) \right\},$$

²There are $n = \sum_{k=0}^{\Gamma} \binom{96}{k}$ possible trajectories where n is within the interval $[2^\Gamma, 2^{96}]$ as $(1+1)^\Gamma = \sum_{k=0}^{\Gamma} \binom{\Gamma}{k}$ and $(1+1)^{96} = \sum_{k=0}^{96} \binom{96}{k}$ by using the binomial formula.

with the set of constraints Φ defined by

$$\phi_t^y - \phi_t^{y^{\min}} + \phi_t^{y^{\max}} - \phi_t^{d^-} + \phi_t^{d^+} = -\pi_t \Delta t, \quad \forall t \in \mathcal{T} \quad [y_t] \quad (12.11a)$$

$$-\phi_t^{d^-} \leq \beta \pi_t \Delta t, \quad \forall t \in \mathcal{T} \quad [d_t^-] \quad (12.11b)$$

$$-\phi_t^{d^+} \leq \beta \pi_t \Delta t, \quad \forall t \in \mathcal{T} \quad [d_t^+] \quad (12.11c)$$

$$\phi_1^{\text{dis}} - \phi_1^y + \phi^{s^{\text{ini}}} \frac{\Delta t}{\eta^{\text{dis}}} \leq 0 \quad [y_1^{\text{dis}}] \quad (12.11d)$$

$$\phi_t^{\text{dis}} - \phi_t^y + \phi_t^s \frac{\Delta t}{\eta^{\text{dis}}} \leq 0, \quad \forall t \in \mathcal{T} \setminus \{1\} \quad [y_t^{\text{dis}}] \quad (12.11e)$$

$$\phi_1^{\text{cha}} + \phi_1^y - \phi^{s^{\text{ini}}} \eta^{\text{cha}} \Delta t \leq 0 \quad [y_1^{\text{cha}}] \quad (12.11f)$$

$$\phi_t^{\text{cha}} + \phi_t^y - \phi_t^s \eta^{\text{cha}} \Delta t \leq 0, \quad \forall t \in \mathcal{T} \setminus \{1\} \quad [y_t^{\text{cha}}] \quad (12.11g)$$

$$-\phi_1^{s^{\min}} + \phi_1^{s^{\max}} + \phi^{s^{\text{ini}}} - \phi_2^s \leq 0 \quad [s_1] \quad (12.11h)$$

$$-\phi_t^{s^{\min}} + \phi_t^{s^{\max}} + \phi_{t-1}^s - \phi_t^s \leq 0, \quad \forall t \in \mathcal{T} \setminus \{1, 2, T\} \quad [s_t] \quad (12.11i)$$

$$-\phi_T^{s^{\min}} + \phi_T^{s^{\max}} + \phi^{s^{\text{end}}} + \phi_T^s \leq 0 \quad [s_T] \quad (12.11j)$$

$$-\phi_t^y + \phi_t^{y^{\text{pv}}} \leq 0, \quad \forall t \in \mathcal{T} \quad [y_t^{\text{pv}}]. \quad (12.11k)$$

Definition 12.1.1 (Sub-Problem). As a result, the worst-case dispatch problem $\max_{u_t \in \mathcal{U}} \left[\min_{y_t \in \Omega(x_t, u_t)} J(x_t, y_t) \right]$ is equivalent to the following problem, which yields the sub-problem (SP) in the Benders-dual cutting plane and column-and-constraint generation algorithms

$$\mathbf{SP} : R(x_t) = \max_{u_t \in \mathcal{U}, \phi_t \in \Phi} G(x_t, u_t, \phi_t). \quad (12.12)$$

Overall, (12.5) becomes a min-max problem

$$\min_{x_t \in \mathcal{X}} \left[\max_{u_t \in \mathcal{U}, \phi_t \in \Phi} G(x_t, u_t, \phi_t) \right], \quad (12.13)$$

that can be solved using a Benders decomposition technique such as BD or CCG, between a master problem, that is linear, and a sub-problem, that is bilinear, since Indeed, G has the following bilinear terms $\phi_t^{y^{\text{pv}}} u_t = \phi_t^{y^{\text{pv}}} \hat{p}_t^{(0.5)} - \phi_t^{y^{\text{pv}}} z_t u_t^{\min}$. It is possible to linearize the products of the binary and continuous variables $z_t \phi_t^{y^{\text{pv}}}$ of G by using a standard integer algebra trick [156] with the following constraints $\forall t \in \mathcal{T}$

$$-M_t^- z_t \leq \alpha_t \leq M_t^+ z_t \quad (12.14a)$$

$$-M_t^-(1 - z_t) \leq \phi_t^{y^{\text{pv}}} - \alpha_t \leq M_t^+(1 - z_t), \quad (12.14b)$$

where M_t^\pm are the big-M's values of $\phi_t^{y^{pv}}$ and α_t is an auxiliary continuous variable. The definition of the uncertainty set (12.8) with binary variables, based on [189, Proposition 2], is essential to linearize G .

12.1.4 Controller formulation

last measured values, and renewable generation intraday point forecasts. It computes at each period t the set-points from t to the last period T of the day. The formulation is the following MILP

$$\min_{y_t \in \Omega(x_t, \hat{y}_t^{pv})} J(x_t, y_t). \quad (12.15)$$

12.2 Solution methodology

Following the methodology described by Bertsimas et al. [18], Jiang et al. [95], a two-level algorithm can be used to solve the min-max (12.5) problem with a Benders-dual cutting plane algorithm.

Definition 12.2.1 (Master Problem). The following master problem (MP) is solved iteratively by adding new constraints to cut off the infeasible or non-optimal solutions and is defined at iteration j by

$$\text{MP : } \min_{x_t \in \mathcal{X}, \theta} \theta \quad (12.16a)$$

$$\theta \geq G(x_t, \alpha_{t,l}, \phi_{t,l}), \quad \forall l \leq j \quad \text{optimality cuts} \quad (12.16b)$$

$$G(x_t, \tilde{\alpha}_{t,k}, \tilde{\phi}_{t,k}) \leq 0, \quad \forall k \leq j, \quad \text{feasibility cuts} \quad (12.16c)$$

where constraints (12.16b) represent the optimality cuts, generated by retrieving the optimal values $\alpha_{t,l}, \phi_{t,l}$ of the SP (12.12), while constraints (12.16c) represent the feasibility cuts, generated by retrieving the extreme rays $\tilde{\alpha}_{t,k}, \tilde{\phi}_{t,k}$ of (12.12), and θ is the optimal value of the second-stage problem.

The MP can compute an optimal solution at iteration j ($x_{t,j}, \theta_j$). Note, that $R(x_{t,j})$ provides an upper bound and θ_{j+1} provides a lower bound to the optimal value of (12.5). Therefore, by iteratively introducing cutting planes (12.16b and (12.16c) from the SP and computing MP (12.16), lower and upper bounds will converge and an optimal solution of (12.5) can be obtained.

12.2.1 Convergence warm start

A four-dimension taxonomy of algorithmic enhancements and acceleration strategies is proposed by [148]: solution generation, solution procedure, decomposition strategy, cut generation. The solution generation is the method used to set trial values for the SP. The quality of these solutions impacts the number of iterations, as the SP uses them to generate cuts and bounds. The standard strategy is to solve the MP without modification. However, heuristics can be used as a warm-start strategy to generate an initial set of tight cuts to strengthen the MP. A simple heuristic is proposed by [116] to generate feasible solutions and a set of good initial cuts. Computational evidence demonstrated the efficiency of this approach in terms of solution quality and time. Therefore, we designed the following warm-start method to improve the Benders convergence by building an initial set of cuts $\{\theta_i\}_{1 \leq i \leq I}$ for the master problem (12.16). It consists of sampling renewable generation trajectories that are assumed to be close to the worst trajectory of \mathcal{U} . Let t_1 and t_f be the time periods corresponding to the first and last non null PV 50% quantile forecast values. If $m = t_f - (t_1 + \Gamma - 1) > 0$, m trajectories are sampled. The m^{th} sampled trajectory is built by setting the Γ values of the PV 50% quantile forecast to the u_t^{\min} lower bound for time periods $t_1 + (m - 1) \leq t \leq t_1 + \Gamma - 1 + (m - 1)$. An additional trajectory is built by setting the Γ maximum values of the PV 50% quantile forecast to u_t^{\min} lower bound. Then, for each sampled trajectory $u_{t,i}$, the MILP formulation (12.3) is used to compute the related engagement plan $x_{t,i}$. Finally, the cut θ_i is built by solving (12.12) where the uncertainty set is a singleton $\mathcal{U} = \{u_{t,i}\}$, and the engagement plan is $x_{t,i}$ to retrieve the optimal values with (12.16b)

$$\theta_i = G(x_{t,i}, \alpha_{t,i}, \phi_{t,i}), \quad i = 1 \dots I.$$

12.2.2 Algorithm convergence

First³, we make the *relatively complete recourse* assumption that the SP is feasible for any engagement plan x_t and generation trajectory u_t . This assumption is valid in the capacity firming framework where curtailment is allowed. If the system faces underproduction where x_t is large, the generation is 0, and the BESS discharged, penalties are applied. If it encounters overproduction where x_t is close to 0, the generation is large, and the BESS is charged, the excess of generation is curtailed. In

³The comments of this subsection apply to the BD and CCG algorithms.

both cases, there is always a feasible dispatch. Notice, when the relatively complete recourse assumption does not hold, [14] propose an extension of the CCG algorithm.

Second, the convergence of the relaxed SP is checked at each iteration of the algorithm by ensuring there is no simultaneous charge and discharge. However, such a situation should not occur because, in case of overproduction, the excess of generation can be curtailed. Simultaneous charging and discharging could indeed be an equivalent solution to dissipate the excess energy. That solution can be avoided in practice by adding a minor penalty for using the storage system. However, we never observed simultaneous charge and discharge over the hundreds of simulations carried out. Thus, it is not required to implement an extension of the BD or CCG algorithm that handles a linear two-stage robust optimization model with a mixed-integer recourse problem such as proposed by [188].

Finally, the overall convergence of the algorithm toward the optimal solution is checked. Indeed, depending on the big-M's values, the algorithm may converge by reducing the gap between the MP and SP. However, it does not ensure an optimal solution. Therefore, once the convergence between the MP and SP is reached at iteration $j = J$, the objective of the MP at J is compared to the objective of the MILP formulation (12.3) using the worst-case generation trajectory $u_t^{*,J}$ as parameters. If the absolute gap $|MILP^J - MP^J|$ is higher than a convergence threshold ϵ , the convergence is not reached. Then, larger big-M's values are set, and the algorithm is restarted until convergence, or a stopping criterion is reached.

12.2.3 Benders-dual cutting plane algorithm

The Benders-dual cutting plane algorithm consists of solving (12.13) without constraints [(10.7a)-(10.7b)] following the procedure previously described, and to obtain a day-ahead robust schedule $x_{t,J}$ at the last iteration J .

Definition 12.2.2 (Benders-dual cutting plane algorithm). The initialization step consists of setting the initial big-M's values $M_t^- = 1$ and $M_t^+ = 0 \forall t \in \mathcal{T}$, the time limit resolution of the sub-problem (12.12) to 10 s, and the threshold convergence ϵ to 0.5 €. Let MP^j , SP^j , be the MP and SP objective values at iteration j , the lower and upper bounds, respectively, and $MILP^J$ the MILP objective value using the worst renewable generation trajectory $u_t^{*,J}$ at iteration J .

Initialization.

Warm-start: build the initial set of cuts $\{\theta_i\}_{1 \leq i \leq I}$.

while $|MILP^J - MP_2^J| > \epsilon$ and $M_t^- < 500$ **do**
 Initialize $j = 0$, solve the MP (12.16) and retrieve $x_{t,0}$.
while the 10 last $|MP^j - SP^j|$ are not $< \epsilon$ **do**
 Solve the SP (12.12) with $x_{t,j}$ as parameters:
if the SP is unbounded **then**
 Retrieve the extreme rays $\tilde{\alpha}_{t,k}, \tilde{\phi}_{t,k}, 0 \leq k \leq j$.
 Add the k -th feasibility cut: $G(x_{t,j}, \tilde{\alpha}_{t,k}, \tilde{\phi}_{t,k}) \leq 0$.
else
 Retrieve the optimal values $\alpha_{t,l}, \phi_{t,l}, 0 \leq l \leq j$.
 Add the l -th optimality cut: $\theta \geq G(x_{t,j}, \alpha_{t,l}, \phi_{t,l})$.
 Update the upper bound $SP^j = R(x_{t,j})$.
 SP check: no simultaneous charge and discharge.
end if
 Solve the MP (12.16): get the optimal values $\theta_j, x_{t,j}$.
 Update the lower bound $MP^j = \theta_j$ and $j = j + 1$.
end while
 $j = J$: convergence between the SP and MP is reached. Check convergence with MILP: get $u_t^{*,J}$ from SP^J and compute $MILP^J$ (12.3).
if $|MILP^J - MP^J| > \epsilon$ **then**
if $M_t^- \leq 50$ **then**
 Update big-M's values $M_t^- = 10 + M_t^- \forall t \in \mathcal{T}$.
else
 Update big-M's values $M_t^- = 100 + M_t^- \forall t \in \mathcal{T}$.
end if
 Reset j to 0 and restart algorithm with a new MP .
end if
end while
 Retrieve the final $x_{t,J}$ engagement.

12.2.4 Column and constraints generation algorithm

We implemented the column and constraints generation procedure proposed by [189, 183]. The following master problem (MP_2) is solved at iteration j

$$\min_{x_t \in \mathcal{X}, \theta, \{y_t^s\}_{0 \leq s \leq j}} \theta \quad (12.17a)$$

$$\theta \geq J(x_t, y_t^s), \quad s = 0 \dots j \quad (12.17b)$$

$$y_t^s \in \Omega(x_t, u_t^{*,s}), \quad s = 0 \dots j, \quad (12.17c)$$

where constraints (12.17b) and (12.17c) serve as optimality and feasibility, respectively. $\{y_t^s\}_{0 \leq s \leq j}$ are the new variables added to the MP_2 , and $u_t^{*,s}$ represent the worst PV trajectory computed by the SP at iteration $0 \leq s \leq j$. Note: in our CCG implementation, we solve the SP with the same approach as the SP of the BD algorithm.

Definition 12.2.3 (Column and constraints generation algorithm). The initialization step is identical to BD, and the CCG algorithm implemented is similar to BD. Note: there is a maximum of 50 iterations between the SP and the MP_2 before checking the convergence with the MILP. If the criterion is not reached, the big-M's values are increased. Indeed, at each iteration j the y_t^j variables are added to the MP_2 . In our case, it represents approximately 1 000 new variables at each iteration. With 50 iterations, the MP_2 is a MILP with approximately 50 000 variables which begins to be hard to solve within a reasonable amount of time.

Initialization.

while $|MILP^j - MP_2^j| > \epsilon$ and $M_t^- < 500$ **do**

Initialize $j = 0$, solve the MP_2 (12.17) and retrieve $x_{t,0}$.

while the two last $|MP_2^j - SP^j|$ are not $< \epsilon$ and $j < 50$ **do**

Solve the SP (12.12) with $x_{t,j}$ as parameters:

Create variables y_t^j in MP_2 .

Retrieve $u_t^{*,j}$ from the SP.

Add the feasibility cut to the MP_2 : $y_t^j \in \Omega(x_t, u_t^{*,j})$.

if the SP is bounded **then**

Add the optimality cut: $\theta \geq J(x_t, y_t^j)$.

Update the upper bound: $SP^j = R(x_{t,j})$.

SP check: no simultaneous charge and discharge.

end if

Solve the MP_2 (12.17): get the optimal values $\theta_j, x_{t,j}$.

Update the lower bound: $MP_2^j = \theta_j$ and $j = j + 1$.

end while

$j = J$: convergence between the SP and MP is reached. Check convergence with MILP: get $u_t^{*,J}$ from SP^J and compute $MILP^J$ (12.3).

if $|MILP^J - MP_2^J| > \epsilon$ **then**

if $M_t^- \leq 50$ **then**

 Update big-M's values $M_t^- = 10 + M_t^- \forall t \in \mathcal{T}$.

else

 Update big-M's values $M_t^- = 100 + M_t^- \forall t \in \mathcal{T}$.

end if

 Reset j to 0 and restart algorithm with a new MP_2 .

end if

end while

Retrieve the final $x_{t,J}$ engagement.

12.3 Case Study

The BD and CCG algorithms are compared on the ULiège case study. It comprises a PV generation plant with an installed capacity $P_c = 466.4$ kWp. The PV generation is monitored on a minute basis, and the data are resampled to 15 minutes. The dataset contains 350 days from August 2019 to November 2020, missing data during March 2020. The NFs approach is compared to a widely used neural architecture, referred to as Long Short-Term Memory (LSTM). In total, eight versions of the planner are considered. Four RO versions: BD-LSTM, BD-NF, CCG-LSTM, and CCG-LSTM. Four deterministic versions: the oracle that uses perfect knowledge of the future, a benchmark that uses PV nominal point forecasts, and two versions using NFs and LSTM PV quantiles. The set of PV quantiles is $\mathcal{Q} = \{q = 10\%, \dots, 50\%\}$. The controller uses PV intraday point forecasts and the day-ahead engagements computed by the planners to compute the set-points and the profits. They are normalized by the profit obtained with the oracle planner and expressed in %.

Section 12.3.1 presents the numerical settings. Section 12.3.2 provides the results of the sensitivity analysis for several risk-averse pairs $[u_t^{min} = \hat{y}_t^{pv,(q)}, \Gamma]$, with $q = 10, \dots, 40\%$, and $\Gamma = 12, 24, 36, 48$. Section 12.3.3 investigates a dynamic risk-averse parameter selection. Section 12.3.4 presents the improvement in terms of computation time provided by the initial set of cuts. Finally, Section 12.3.5 compares the BD and CCG algorithms.

12.3.1 Numerical settings

The testing set is composed of thirty days randomly selected from the dataset. The simulation parameters of the planners and the controller are identical. The planning and controlling periods duration are $\Delta t = 15$ minutes. The peak hours are set between 7 pm and 9 pm (UTC+0). The ramping power constraint on the engagements are $\Delta X_t = 7.5\%P_c$ ($15\%P_c$) during off-peak (peak) hours. The lower bounds on the engagement x^{\min} and the net power y^{\min} are set to 0 kW. The upper bound on the engagement x^{\max} and the net power y^{\max} are set to P_c . Finally, the engagement tolerance is $pP_c = 1\%P_c$, and the penalty factor $\beta = 5$. The BESS minimum s^{\min} and maximum s^{\max} capacity are 0 kWh and 466.4 kWh, respectively. It is assumed to be capable of fully charging or discharging in one hour $y_{\max}^{\text{dis}} = y_{\max}^{\text{cha}} = s^{\max}/1$ with charging and discharging efficiencies $\eta^{\text{cha}} = \eta^{\text{dis}} = 95\%$. Each simulation day is independent with a fully discharged battery at the first and last period $s^{\text{ini}} = s^{\text{end}} = 0$ kWh. The Python Gurobi library is used to implement the algorithms in Python 3.7, and Gurobi⁴ 9.0.2 to solve all the optimization problems. Numerical experiments are performed on an Intel Core i7-8700 3.20 GHz based computer with 12 threads and 32 GB of RAM running on Ubuntu 18.04 LTS.

Figures 12.3a and 12.3b illustrate the LSTM and NFs PV quantile forecasts, observation, and nominal point forecasts on *September 14*, 2019. Figures 12.3c and 12.3d provide the engagement plan (x) and the BESS state of charge (s) computed with the RO planner, the deterministic planner with the nominal point forecasts, and the perfect knowledge of the future.

12.3.2 Constant risk-averse parameters strategy

The risk-averse parameters of the RO approach $[\hat{y}_t^{\text{pv},(q)}, \Gamma]$ are constant over the dataset. One way to identify the optimal pair is to perform a sensitivity analysis [174]. Figure 12.4 provides the normalized profits of the BD-RO, CCG-RO, and deterministic planners using PV quantiles, left with LSTM and right with NFs, and nominal point forecasts. The RO and deterministic planners outperform by a large margin the baseline. The latter, the deterministic planner with nominal point forecasts, cannot deal with PV uncertainty and achieved only 53.3 %. Then, the planners using NFs quantiles significantly outperform the planners with LSTM quantiles. Overall, the CCG algorithm achieved better results for almost all pairs of risk-averse parameters. The highest profits achieved by the CCG-NF, BD-NF and NF-deterministic planners are 73.8

⁴<https://www.gurobi.com/>

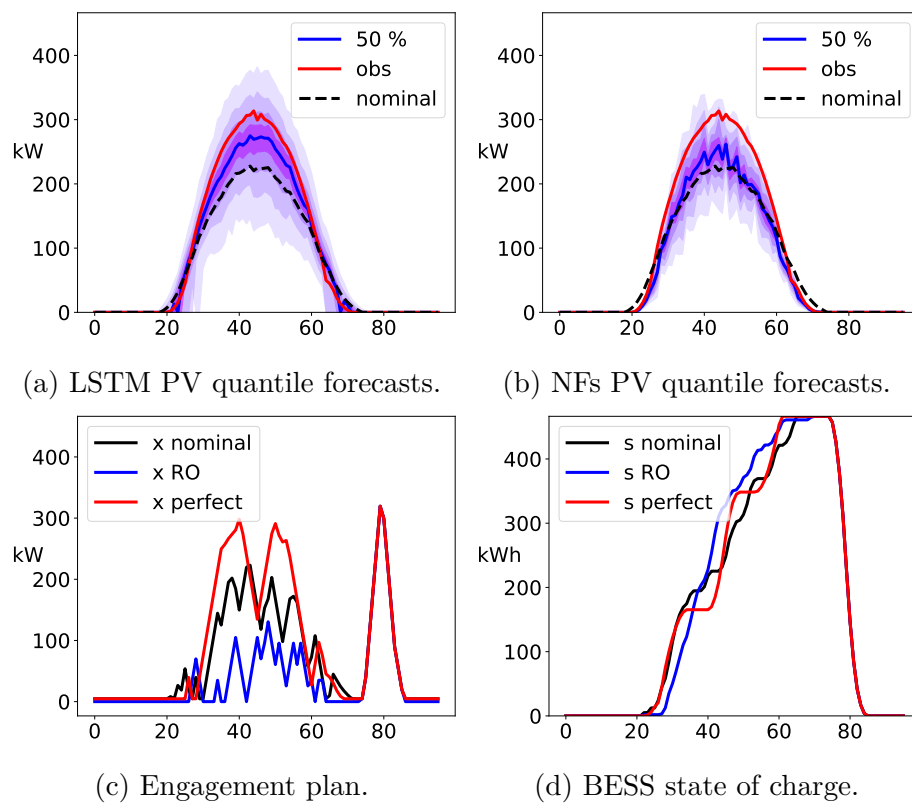


Fig. 12.3 Results illustration on *September 14, 2019*.

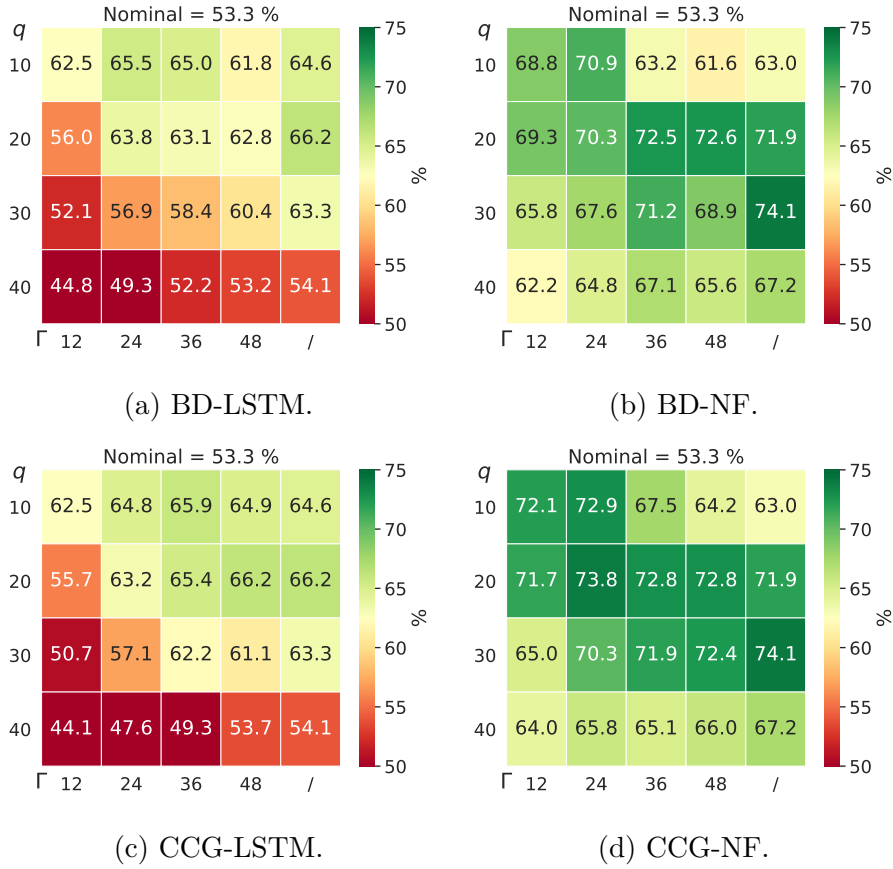


Fig. 12.4 Results with constant risk-averse parameters. Normalized profit (%) of the BD and CCG RO planners $([\Gamma, q])$, deterministic $([/, q])$ planner, and the reference that is the deterministic planner with point-forecasts (Nominal). Left part: LSTM quantiles, right part: NF quantiles.

%, 72.6 % and 74.1 %, respectively, with the risk-averse parameters $[q = 20\%, \Gamma = 24]$, $[q = 20\%, \Gamma = 48]$, and the quantile 30 %. It should be possible to improve the RO results by tuning the risk-averse parameters $[\hat{y}_t^{\text{pv},(q)}, \Gamma]$. However, these results emphasize the interest in considering a deterministic planner with the relevant PV quantile as point forecasts, which are easy to implement, fast to compute (a few seconds), and less prone to convergence issues than the RO approach.

12.3.3 Dynamic risk-averse parameters strategy

In this section, the risk-averse parameters $[u_t^{\text{min}}, \Gamma]$ of the RO approach are dynamically set based on the day-ahead quantile forecasts distribution, and u_t^{min} is not necessarily equal to the same quantile $\hat{y}_t^{\text{pv},(q)} \forall t \in \mathcal{T}$. The motivation of this strategy is to assume

that the sharper the quantile forecast distribution around the median is, the less risk-averse the RO approach should be.

Two parameters are designed to this end: (1) the PV uncertainty set max depth d_q to control u_t^{min} ; (2) the budget depth d_Γ to control Γ . d_q is a percentage of the distance between the median and the 10% quantile d_{50-10} , and d_Γ is a percentage of the total installed capacity P_c . Then, two rules are designed to dynamically set the risk-averse parameters $[u_t^{min}, \Gamma]$ for each day of the dataset. For a given day, and the set of time periods where the PV median is non null, the distances between the PV median and the PV quantiles 20, 30, and 40% are computed: d_{50-20} , d_{50-30} , d_{50-40} . u_t^{min} is dynamically set at each time period t as follows

$$u_t^{min} = \begin{cases} \hat{y}_t^{pv,(0.1)} & \text{if } d_t^{50-20/30/40} > d_q d_t^{50-10} \\ \hat{y}_t^{pv,(0.2)} & \text{if } d_t^{50-20/30} > d_q d_t^{50-10} \\ \hat{y}_t^{pv,(0.3)} & \text{if } d_t^{50-20} > d_q d_t^{50-10} \\ \hat{y}_t^{pv,(0.4)} & \text{otherwise} \end{cases}. \quad (12.18)$$

For a given day, the budget of uncertainty Γ is dynamically set based on the following rule

$$\Gamma = \#\{t : d_t^{50-10} > d_\Gamma P_c\}. \quad (12.19)$$

Figure 12.5 provides the normalized profits of the CCG-RO, BD-RO, and deterministic planners for several pairs $[d_\Gamma, d_q]$ using both the LSTM and NF quantiles. The planners achieved better results when using the NF quantiles. Overall, the results are improved compared to fixed risk-averse parameters for all the planners. The highest profits achieved by the CCG-NF, BD-NF and NF-deterministic planners are 75.0 %, 72.6 % and 75.0 %, respectively, with $[d_\Gamma, d_q] = [10, 30]$, $[d_\Gamma, d_q] = [10, 5]$, and $d_q = 50\%$.




Fig. 12.5 Results with dynamic risk-averse parameters. Normalized profit (%) of the BD and CCG RO planners ($[d_r, d_q]$), and deterministic ($[/, d_q]$) planner. Left part: LSTM quantiles, right part: NF quantiles.

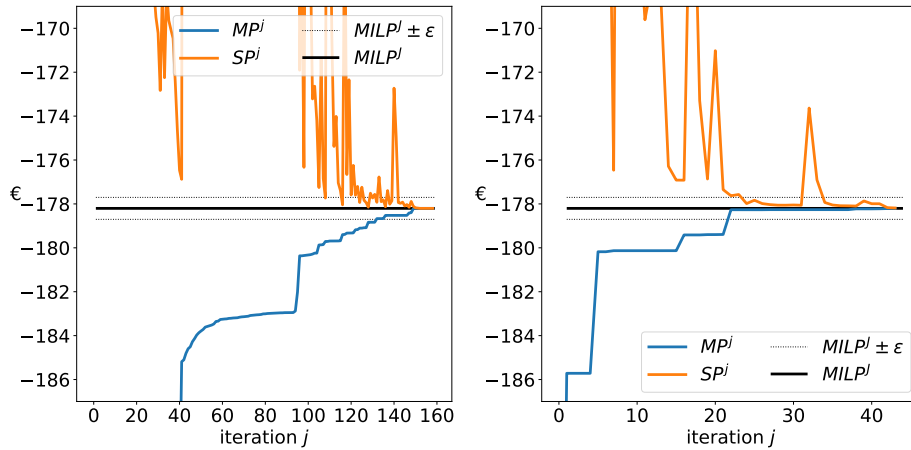


Fig. 12.6 Benders convergence without (left) and with (right) an initial set of cuts on *September 14, 2019*.

$\{\theta_i\}_{1 \leq i \leq I}$	t^{av}	$t^{50\%}$	t^{min}	t^{max}	t^{tot}
False	3.5	2.0	< 0.1	34.1	105.4
True	2.0	0.7	< 0.1	30.4	61.3

Table 12.1 Computation times (min) statistics.

12.3.4 BD convergence warm start improvement

Overall, the warm-start procedure of the BD algorithm improves the convergence by reducing the number of iterations and allows to reach more frequently optimal solution. In addition, it reduces the number of times the big-M's values need to be increased before reaching the final convergence criterion with the MILP. It is illustrated by considering the dynamic risk-averse parameters strategy with $[d_\Gamma, d_q] = [10, 10]$. Figure 12.6 illustrates the reduction of the total number of iteration J required to converge below the threshold ϵ , on a specific day of the dataset. It is divided by 3.6 from 159 to 44. The computation time is divided by 4.1 from 7.4 min to 1.8 min. Table 12.1 provides the computation times (min) statistics over the entire dataset with and without warm start. The averaged t^{av} and total t^{tot} computation times are drastically reduced when using the warm-start.

Algorithm	RO-type	\bar{t}	1%	J^{max}
BD-NF	static	85.2 (151.9)	0.0	72.6
CCG-NF	static	7.5 (6.0)	1.9	73.8
BD-NF	dynamic	102.3 (107.3)	0.0	72.6
CCG-NF	dynamic	9.2 (5.5)	4.2	75.0

Table 12.2 BD vs CCG statistics.

\bar{t} (s) is the averaged computation time per day with the standard deviation in the bracket. 1% (%) is the % of instances that did not terminate with optimality. \bar{t} and 1% are computed over all days of the testing set and for all pair of constant (dynamic) risk-averse parameters $[u_t^{min}, \Gamma]$ ($[d_\Gamma, d_q]$). J^{max} (%) is the best-normalized profit achieved using the NF quantiles over all risk-averse parameters.

12.3.5 BD and CCG comparison

Table 12.2 provides a comparison of the BD and CCG algorithms when using NF quantiles for both the static and dynamic robust optimization strategies. Overall, the CCG algorithm converges in 5-10 iterations instead of 50-100 for BD. Therefore, the CCG computes the day-ahead planning in approximately 10 seconds, ten times faster than BD. This observation is consistent with [183] that demonstrated the CCG algorithm converges faster than BD. Let n be the number of extreme points of the uncertainty set \mathcal{P} and m of the space Φ defined by constraint (12.11). The BD algorithm computes an optimal solution in $O(nm)$ iterations, and the CCG procedure in $O(n)$ iterations [183]. Note: the BD algorithm is still competitive in an operational framework as it takes on average 1-2 minutes to compute the day-ahead planning. However, we observed that the CCG does not always converge to an optimal solution (see Section 12.2.2), which never happened with the BD algorithm. Fortunately, these cases amount to only a few % of the total instances. Overall, the CCG algorithm achieved better results than the BD for almost all the risk-averse parameters. Finally, in our opinion, both algorithms require the same amount of knowledge to be implemented. Indeed, the only difference is the MP as the SP are solved identically.

12.4 Conclusion

The core contribution of this study is to address the two-phase engagement/control problem in the context of capacity firming. A secondary contribution is to use a recent deep learning technique, Normalizing Flows, to compute PV quantiles. It is compared to a typical neural architecture, referred to as Long Short-Term Memory. We developed an integrated forecast-driven strategy modeled as a min-max-min robust optimization problem with recourse that is solved using a Benders decomposition procedure. Two main cutting plane algorithms used to address the two-stage RO unit commitment problems are compared: the Benders-dual cutting plane and the column-and-constraint generation algorithms. The convergence is checked by ensuring a gap below a threshold between the final objective and the corresponding deterministic, objective value. A risk-averse parameter assessment selects the optimal robust parameters and the optimal conservative quantile for the deterministic planner. Both the NF-based and LSTM-based planners outperformed the deterministic planner with nominal point PV forecasts. The NF model outperforms the LSTM in forecast value as the planner using the NF quantiles achieved higher profit than the planner with LSTM quantiles. Finally, a dynamic risk-averse parameter selection strategy is built by taking advantage of the PV quantile forecast distribution and provides further improvements. The CCG procedure converges ten times faster than the BD algorithm in this case study and achieves better results. However, it does not always converge to an optimal solution.

Overall, the RO approach for both the BD and CCG algorithms allows finding a trade-off between conservative and risk-seeking policies by selecting the optimal robust optimization parameters, leading to improved economic benefits compared to the baseline. Therefore, offering a probabilistic guarantee for the robust solution. However, the deterministic planner with the relevant PV quantile achieved interesting results. It emphasizes the interest to consider a well-calibrated deterministic approach. Indeed, it is easy to implement, computationally tractable for large-scale problems, and less prone to convergence issues. Note: this approach can be used in any other case study. It only requires a few months of data, renewable generation, and weather forecasts to train the forecasting models to compute reliable forecasts for the planner.

Several extensions are under investigation: (1) a stochastic formulation of the planner with improved PV scenarios based on Gaussian copula methodology or generated by a state-of-the-art deep learning technique such as Normalizing Flows, Generative Adversarial Networks or Variational AutoEncoders; (2) an improved dynamic risk-averse parameter selection strategy based on a machine learning tool capable of better-taking advantage of the PV quantiles distribution.

Chapter 13

Energy retailer

Overview

This Chapter investigates the forecast value assessment of the deep generative models studied in Chapter 6. The reader is referred to Chapter 6 for the context, primary contributions, the NFs, GANs, and VAEs background, and the description of the Global Energy Forecasting Competition 2014 (GEFcom 2014) case study.

References: This chapter is an adapted version of the following publication:

Jonathan Dumas, Antoine Wehenkel, Damien Lanaspeze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.

“*Be willing to make decisions. That’s the
most important quality in a good leader.*”

— **George S. Patton**

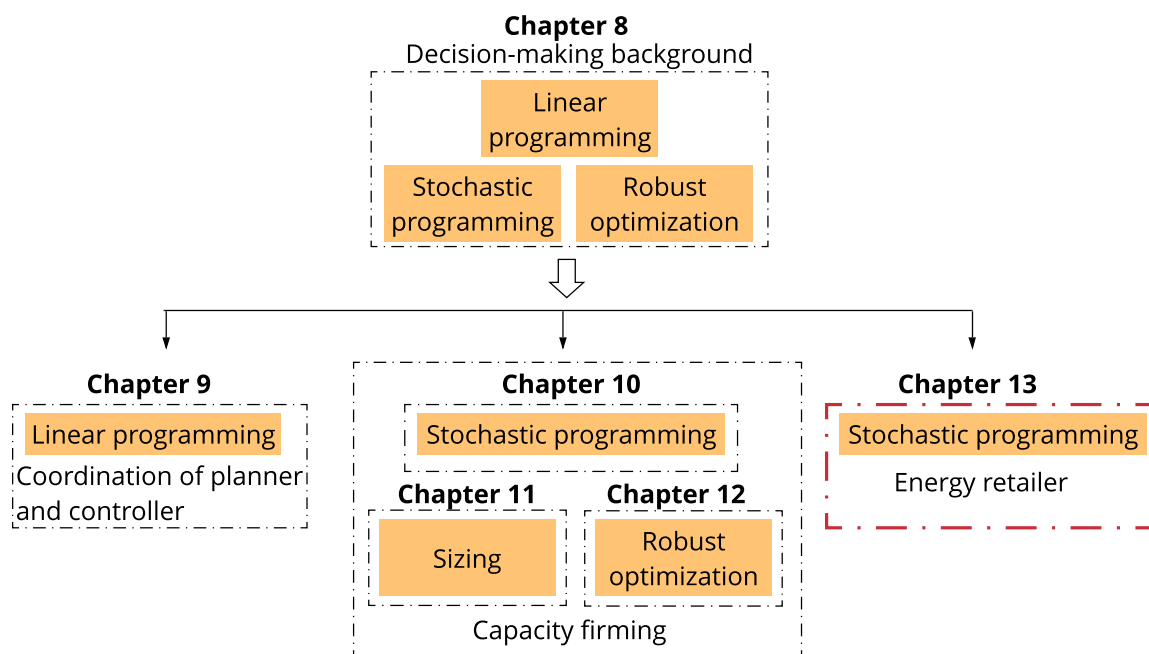


Fig. 13.1 Chapter 13 position in Part II.

A model that yields lower errors in terms of forecast quality may not always point to a more effective model for forecast practitioners [87]. To this end, similarly to Toubeau et al. [166], the forecast value is assessed by considering the day-ahead market scheduling of electricity aggregators, such as energy retailers or generation companies. The energy retailer aims to balance its portfolio on an hourly basis to avoid financial penalties in case of imbalance by exchanging the surplus or deficit of energy in the day-ahead electricity market. The energy retailer may have a battery energy storage system (BESS) to manage its portfolio and minimize imports from the main grid when day-ahead prices are prohibitive.

Section 13.2 introduces the notations used in this Chapter. Section 13.2 presents the formulation of the energy retailer case study. Section 13.3 details empirical value results on the GEFcom 2014 dataset, and Section 13.4 summarizes the main findings and highlights ideas for further work.

13.1 Notation

Sets and indices

Name	Description
t	Time period index.
ω	Scenario index.
T	Number of time periods per day.
$\#\Omega$	Number of scenarios.
\mathcal{T}	Set of time periods, $\mathcal{T} = \{1, 2, \dots, T\}$.
Ω	Set of scenarios, $\Omega = \{1, 2, \dots, \#\Omega\}$.

Parameters

Name	Description	Unit
e_t^{min}, e_t^{max}	Minimum/maximum day-ahead bid.	MWh
y_t^{min}, y_t^{max}	Minimum/maximum retailer net position.	MWh
$y_{max}^{dis}, y_{max}^{cha}$	BESS maximum (dis)charging power.	MW
η^{dis}, η^{cha}	BESS (dis)charging efficiency.	-
s^{min}, s^{max}	BESS minimum/maximum capacity.	MWh
s^{ini}, s^{end}	BESS initial/final state of charge.	MWh
π_t	Day-ahead price.	€/MWh

$\bar{q}_t, \bar{\lambda}_t$	Negative/positive imbalance price.	€/MWh
Δt	Duration of a time period.	hour

Variables

For the sake of clarity the subscript ω is omitted.

Name	Range	Description	Unit
e_t	$[e_t^{\min}, e_t^{\max}]$	Day-ahead bid.	MWh
y_t	$[y_t^{\min}, y_t^{\max}]$	Retailer net position.	MWh
y_t^{PV}	$[0, 1]$	PV generation.	MW
y_t^{W}	$[0, 1]$	Wind generation.	MW
y_t^{l}	$[0, 1]$	Load.	MW
y_t^{cha}	$[0, y_{\text{max}}^{\text{cha}}]$	Charging power.	MW
y_t^{dis}	$[0, y_{\text{max}}^{\text{dis}}]$	Discharging power.	MW
s_t	$[s^{\min}, s^{\max}]$	BESS state of charge.	MWh
d_t^-, d_t^+	\mathbb{R}_+	Short/long deviation.	MWh
y_t^b	$\{0, 1\}$	BESS binary variable.	-

13.2 Problem formulation

Let e_t [MWh] be the net energy retailer position on the day-ahead market during the t -th hour of the day, modeled as a first stage variable. Let y_t [MWh] be the realized net energy retailer position during the t -th hour of the day, which is modeled as a second stage variable due to the stochastic processes of the PV generation, wind generation, and load. Let π_t [€/ MWh] the clearing price in the spot day-ahead market for the t -th hour of the day, q_t ex-post settlement price for negative imbalance $y_t < e_t$, and λ_t ex-post settlement price for positive imbalance $y_t > e_t$. The energy retailer is assumed to be a price taker in the day-ahead market. It is motivated by the individual energy retailer capacity being negligible relative to the whole market. The forward settlement price π_t is assumed to be fixed and known. As imbalance prices tend to exhibit volatility and are difficult to forecast, they are modeled as random variables, with expectations denoted by $\bar{q}_t = \mathbb{E}[q_t]$ and $\bar{\lambda}_t = \mathbb{E}[\lambda_t]$. They are assumed to be independent random variables from the energy retailer portfolio.

13.2.1 Stochastic planner

A stochastic planner with a linear programming formulation and linear constraints is implemented using a scenario-based approach. The planner computes the day-

ahead bids e_t that cannot be modified in the future when the uncertainty is resolved. The second stage corresponds to the dispatch decisions $y_{t,\omega}$ in scenario ω that aims at avoiding portfolio imbalances modeled by a cost function f^c . The second-stage decisions are therefore scenario-dependent and can be adjusted according to the realization of the stochastic parameters. The stochastic planner objective to maximize is

$$J_S = \mathbb{E} \left[\sum_{t \in \mathcal{T}} \pi_t e_t + f^c(e_t, y_{t,\omega}) \right], \quad (13.1)$$

where the expectation is taken with respect to the random variables, the PV generation, wind generation, and load. Using a scenario-based approach, (13.1) is approximated by

$$J_S \approx \sum_{\omega \in \Omega} \alpha_\omega \sum_{t \in \mathcal{T}} \left[\pi_t e_t + f^c(e_t, y_{t,\omega}) \right], \quad (13.2)$$

with α_ω the probability of scenario $\omega \in \Omega$, and $\sum_{\omega \in \Omega} \alpha_\omega = 1$. The mixed-integer linear programming (MILP) optimization problem to solve is

$$\max_{e_t \in \mathcal{X}, y_{t,\omega} \in \mathcal{Y}(e_t)} \sum_{\omega \in \Omega} \alpha_\omega \sum_{t \in \mathcal{T}} \left[\pi_t e_t - \bar{q}_t d_{t,\omega}^- - \bar{\lambda}_t d_{t,\omega}^+ \right], \quad (13.3a)$$

$$\mathcal{X} = \left\{ e_t : e_t \in [e_t^{\min}, e_t^{\max}] \right\}, \quad (13.3b)$$

$$\mathcal{Y}(e_t) = \left\{ y_{t,\omega} : (13.4a) - (13.4m) \right\}. \quad (13.3c)$$

The optimization variables are e_t , day-ahead bid of the net position, $\forall \omega \in \Omega$, $y_{t,\omega}$, retailer net position in scenario ω , $d_{t,\omega}^-$, short deviation, $d_{t,\omega}^+$, long deviation, $y_{t,\omega}^{\text{PV}}$, PV generation, $y_{t,\omega}^{\text{w}}$, wind generation, $y_{t,\omega}^{\text{cha}}$, battery energy storage system (BESS) charging power, $y_{t,\omega}^{\text{dis}}$, BESS discharging power, $s_{t,\omega}$, BESS state of charge, and $y_{t,\omega}^b$ a binary variable to prevent from charging and discharging simultaneously. The imbalance penalty is modeled by the constraints (13.4a)-(13.4b) $\forall \omega \in \Omega$, that define the short and long deviations variables $d_{t,\omega}^-, d_{t,\omega}^+ \in \mathbb{R}_+$. The energy balance is provided by (13.4c) $\forall \omega \in \Omega$. The set of constraints that bound $y_{t,\omega}^{\text{PV}}$ and $y_{t,\omega}^{\text{w}}$ variables are (13.4d)-(13.4e) $\forall \omega \in \Omega$ where $\hat{y}_{t,\omega}^{\text{PV}}$ and $\hat{y}_{t,\omega}^{\text{w}}$ are PV and wind generation scenarios. The load is assumed to be non-flexible and is a parameter (13.4f) $\forall \omega \in \Omega$ where $\hat{y}_{t,\omega}^{\text{l}}$ are load scenarios. The BESS constraints are provided by (13.4g)-(13.4j), and the BESS dynamics by

(13.4k)-(13.4m) $\forall \omega \in \Omega$.

$$-d_{t,\omega}^- \leq -(e_t - y_{t,\omega}), \forall t \in \mathcal{T} \quad (13.4a)$$

$$-d_{t,\omega}^+ \leq -(y_{t,\omega} - e_t), \forall t \in \mathcal{T} \quad (13.4b)$$

$$\frac{y_{t,\omega}}{\Delta t} = y_{t,\omega}^{\text{pv}} + y_{t,\omega}^{\text{w}} - y_{t,\omega}^{\text{l}} + y_{t,\omega}^{\text{dis}} - y_{t,\omega}^{\text{cha}}, \forall t \in \mathcal{T} \quad (13.4c)$$

$$y_{t,\omega}^{\text{pv}} \leq \hat{y}_{t,\omega}^{\text{pv}}, \forall t \in \mathcal{T} \quad (13.4d)$$

$$y_{t,\omega}^{\text{w}} \leq \hat{y}_{t,\omega}^{\text{w}}, \forall t \in \mathcal{T} \quad (13.4e)$$

$$y_{t,\omega}^{\text{l}} = \hat{y}_{t,\omega}^{\text{l}}, \forall t \in \mathcal{T} \quad (13.4f)$$

$$y_{t,\omega}^{\text{cha}} \leq y_{t,\omega}^{\text{b}} y_{\text{max}}^{\text{cha}}, \forall t \in \mathcal{T} \quad (13.4g)$$

$$y_{t,\omega}^{\text{dis}} \leq (1 - y_{t,\omega}^{\text{b}}) y_{\text{max}}^{\text{dis}}, \forall t \in \mathcal{T} \quad (13.4h)$$

$$-s_{t,\omega} \leq -s^{\text{min}}, \forall t \in \mathcal{T} \quad (13.4i)$$

$$s_{t,\omega} \leq s^{\text{max}}, \forall t \in \mathcal{T} \quad (13.4j)$$

$$\frac{s_{1,\omega} - s^{\text{ini}}}{\Delta t} = \eta^{\text{cha}} y_{1,\omega}^{\text{cha}} - \frac{y_{1,\omega}^{\text{dis}}}{\eta^{\text{dis}}}, \quad (13.4k)$$

$$\frac{s_{t,\omega} - s_{t-1,\omega}}{\Delta t} = \eta^{\text{cha}} y_{t,\omega}^{\text{cha}} - \frac{y_{t,\omega}^{\text{dis}}}{\eta^{\text{dis}}}, \forall t \in \mathcal{T} \setminus \{1\} \quad (13.4l)$$

$$s_{T,\omega} = s^{\text{end}} = s^{\text{ini}}. \quad (13.4m)$$

Notice that if $\bar{\lambda}_t < 0$, the surplus quantity is remunerated with a non-negative price. In practice, such a scenario could be avoided provided that the energy retailer has curtailment capabilities, and $(\bar{q}_t, \bar{\lambda}_t)$ are strictly positive in our case study. The deterministic formulation with perfect forecasts, the oracle (O), is a specific case of the stochastic formulation by considering only one scenario where $y_{t,\omega}^{\text{pv}}$, $y_{t,\omega}^{\text{w}}$, and $y_{t,\omega}^{\text{l}}$ become the actual values of PV, wind, and load $\forall t \in \mathcal{T}$. The optimization variables are e_t , y_t , d_t^- , d_t^+ , y_t^{pv} , and y_t^{w} , y_t^{cha} , y_t^{dis} , s_t , and y_t^{b} .

13.2.2 Dispatching

Once the bids e_t have been computed by the planner, the dispatching consists of computing the second stage variables given observations of the PV, wind power, and load. The dispatch formulation is a specific case of the stochastic formulation with e_t as parameter and by considering only one scenario where $y_{t,\omega}^{\text{pv}}$, $y_{t,\omega}^{\text{w}}$, and $y_{t,\omega}^{\text{l}}$ become the actual values of PV, wind, and load $\forall t \in \mathcal{T}$. The optimization variables are y_t , d_t^- , d_t^+ , y_t^{pv} , and y_t^{w} , y_t^{cha} , y_t^{dis} , s_t , and y_t^{b} .




Fig. 13.2 Energy retailer case study: illustration of the observations on a random day of the testing set.

The energy retailer portfolio comprises PV generation, wind power, load, and storage device. The PV, wind power, and load scenarios from the testing set are used as inputs of the stochastic day-ahead planner to compute the optimal bids. The net is the power balance of the energy retailer portfolio. The day-ahead prices π_t are obtained from the Belgian day-ahead market on *February 6, 2020*.

13.3 Value results

The energy retailer portfolio comprises wind power, PV generation, load, and a battery energy storage device. The 50 days of the testing set are used and combined with the 30 possible PV and wind generation zones (three PV zones and ten wind farms), resulting in 1 500 independent simulated days. A two-step approach is employed to evaluate the forecast value:

- First, for each generative model and the 1 500 days simulated, the two-stage stochastic planner computes the day-ahead bids of the energy retailer portfolio using the PV, wind power, and load scenarios. After solving the optimization, the day-ahead decisions are recorded.
- Then, a real-time dispatch is carried out using the PV, wind power, and load observations, with the day-ahead decisions as parameters.

This two-step methodology is applied to evaluate the three generative models, namely the NF, GAN, and VAE. Figure 13.2 illustrates an arbitrary random day of the testing set with the first zone for both the PV and wind. π_t [€/ MWh] is the day-ahead prices on *February 6, 2020* of the Belgian day-ahead market used for the 1 500 days simulated. The negative \bar{q}_t and positive $\bar{\lambda}_t$ imbalance prices are set to $2 \times \pi_t, \forall t \in \mathcal{T}$. The retailer aims to balance the net power, the red curve in Figure 13.2, by importing/exporting from/to the main grid. Usually, the net is positive (negative) at noon (evening) when

the PV generation is maximal (minimal), and the load is minimal (maximal). As the day-ahead spot price is often maximal during the evening load peak, the retailer seeks to save power during the day by charging the battery to decrease the import during the evening. Therefore, the more accurate the PV, wind generation, and load scenarios are, the better is the day-ahead planning.

The battery minimum s^{\min} and maximum s^{\max} capacities are 0 and 1, respectively. It is assumed to be capable of fully (dis)charging in two hours with $y_{\max}^{\text{dis}} = y_{\max}^{\text{cha}} = s^{\max}/2$, and the (dis)charging efficiencies are $\eta^{\text{dis}} = \eta^{\text{cha}} = 95\%$. Each simulation day is independent with a fully discharged battery at the first and last period of each day $s^{\text{ini}} = s^{\text{end}} = 0$. The 1 500 stochastic optimization problems are solved with 50 PV, wind generation, and load scenarios. The python Gurobi library is used to implement the algorithms in Python 3.7, and Gurobi¹ 9.0.2 is used to solve the optimization problems. Numerical experiments are performed on an Intel Core i7-8700 3.20 GHz based computer with 12 threads and 32 GB of RAM running on Ubuntu 18.04 LTS.

The net profit, that is, the profit minus penalty, is computed for the 1 500 days of the simulation and aggregated in the first row of Table 13.1. The ranking of each model is computed for the 1 500 days. The cumulative ranking is expressed in terms of percentage in Table 13.1. NF outperformed both the GAN and VAE with a total net profit of 107 k€. There is still room for improvement as the oracle, which has perfect future knowledge, achieved 300 k€. NF, ranked first 39.0% during the 1 500 simulation days and achieved the first and second ranks 69.6%. Overall, in terms of

	NF	VAE	GAN
Net profit (k€)	107	97	93
1 (%)	39.0	31.8	29.2
1 & 2 (%)	69.6	68.3	62.1
1 & 2 & 3 (%)	100	100	100

Table 13.1 Total net profit (k€) and cumulative ranking (%).

Using the NF PV, wind power, and load scenarios, the stochastic planner achieved the highest net profit with 107 k€, ranked first 39.0 %, second 30.6 %, and third 30.4 % over 1 500 days of simulation. The second-best model, the VAE, achieved a net profit of 97 k€, ranked first 31.8 %, second 36,5 %, and third 31.7 %.

forecast value, the NF outperforms the VAE and GAN. However, this case study is "simple," and stochastic optimization relies mainly on the quality of the average of

¹<https://www.gurobi.com/>

the scenarios. Therefore, one may consider taking advantage of the particularities of a specific method by considering more advanced case studies. In particular, the specificity of the NFs to provide direct access to the probability density function may be of great interest in specific applications. It is left for future investigations as more advanced case studies would prevent a fair comparison between models.

13.3.1 Results summary

Table 13.2 summarizes the main results of this study by comparing the VAE, GAN, and NF implemented through easily comparable star ratings. The rating for each criterion is determined using the following rules - 1 star: third rank, 2 stars: second rank, and 3 stars: first rank. Specifically, training speed is assessed based on reported total training times for each dataset: PV generation, wind power, and load; sample speed is based on reported total generating times for each dataset; quality is evaluated with the metrics considered; value is based on the case study of the day-ahead bidding of the energy retailer; the hyper-parameters search is assessed by the number of configurations tested before reaching satisfactory and stable results over the validation set; the hyper-parameters sensitivity is evaluated by the impact on the quality metric of deviations from the optimal the hyper-parameter values found during the hyper-parameter search; the implementation-friendly criterion is appraised regarding the complexity of the technique and the amount of knowledge required to implement it.

13.4 Conclusions

This Chapter proposed a fair and thorough comparison of NFs with the state-of-the-art deep learning generative models, GANs and VAEs, in terms of value. The experiments are performed using the open data of the Global Energy Forecasting Competition 2014. The generative models use the conditional information to compute improved weather-based PV, wind power, and load scenarios. This study demonstrated that NFs are capable of challenging GANs and VAEs. They are overall more accurate both in terms of quality (see Chapter 6) and value and can be used effectively by non-expert deep learning practitioners. Note, Section 6.6 in Chapter 6 presents the NFs advantages over more traditional deep learning approaches that should motivate their introduction into power system applications.

Criteria	VAE	GAN	NF
Train speed	***	***	***
Sample speed	***	***	***
Quality	***	***	***
Value	***	***	***
Hp search	***	***	***
Hp sensibility	***	***	***
Implementation	***	***	***

Table 13.2 Comparison between the deep generative models.

The rating for each criterion is determined using the following rules - 1 star: third rank, 2 stars: second rank, and 3 stars: first rank. Train speed: training computation time; Sample speed: scenario generation computation time; Quality: forecast quality based on the eight complementary metrics considered; Value: forecast value based on the day-ahead energy retailer case study; Hp search: assess the difficulty to identify relevant hyper-parameters; Hp sensibility: assess the sensitivity of the model to a given set of hyper-parameters (the more stars, the more robust to hyper-parameter modifications); Implementation: assess the difficulty to implement the model (the more stars, the more implementation-friendly). Note: the justifications are provided in Appendix 13.5.

13.5 Appendix: Table 13.2 justifications

The VAE is the fastest to train, with a recorded computation time of 7 seconds on average per dataset. The training time of the GAN is approximately three times longer, with an average computation time of 20 seconds per dataset. Finally, the NF is the slowest, with an average training time of 4 minutes. This ranking is preserved with the VAE the fastest concerning the sample speed, followed by the GAN and NF models. The VAE and the GAN generate the samples over the testing sets, 5 000 in total, in less than a second. However, the NF considered takes a few minutes. In contrast, the affine autoregressive version of the NF is much faster to train and generate samples. Note: even a training time of a few hours is compatible with day-ahead planning applications. In addition, once the model is trained, it is not necessarily required to retrain it every day.

The quality and value assessments have already been discussed in Sections 6.5 and 13.3. Overall, the NF outperforms both the VAE and GAN models.

Concerning the hyper-parameters search and sensibility, the NF tends to be the most straightforward model to calibrate. Compared with the VAE and GAN, we found relevant hyper-parameter values by testing only a few combinations. In addition, the NF is robust to hyper-parameter modifications. In contrast, the GAN is the most sensitive. Variations of the hyper-parameters may result in very poor scenarios both in terms of quality and shape. Even for a fixed set of hyper-parameters values, two separate training may not converge towards the same results illustrating the GAN training instabilities. The VAE is more accessible to train than the GAN but is also sensitive to hyper-parameters values. However, it is less evident than the GAN.

Finally, we discuss the implementation-friendly criterion of the models. Note: this discussion is only valid for the models implemented in this study. There exist various architectures of GANs, VAEs, and NFs with simple and complex versions. In our opinion, the VAE is the effortless model to implement as the encoder and decoder are both simple feed-forward neural networks. The only difficulty lies in the reparameterization trick that should be carefully addressed. The GAN is a bit more difficult to deploy due to the gradient penalty to handle but is similar to the VAE with both the discriminator and the generator that are feed-forward neural networks. The NF is the most challenging model to implement from scratch because the UMNN-MAF approach requires an additional integrand network. An affine autoregressive NF is easier to implement. However, it may be less capable of modeling the stochasticity of the variable of interest. However, forecasting practitioners do not necessarily have

to implement generative models from scratch and can use numerous existing Python libraries.

Chapter 14

Part II conclusions

“Everything has been figured out, except how to live.”

— **Jean-Paul Sartre**

Part II presents several approaches to address the energy management of a grid-connected microgrid. It proposes handling the renewable generation uncertainty by considering deterministic, stochastic, and robust approaches that rely on point forecasts, scenarios, and quantiles forecasts. Several numerical experiments evaluate these approaches.

- The MiRIS microgrid located at the John Cockerill Group’s international headquarters in Seraing, Belgium, assesses a value function-based approach to propagate information from operational planning to real-time optimization. This methodology relies on point forecasts of PV and electrical demand to conduct the deterministic optimization. The results demonstrate the efficiency of this method to manage the peak in comparison with a Rule-Based Controller.
- The MiRIS case study is also employed in the capacity firming market. It allows studying a stochastic approach to address the energy management of a grid-connected renewable generation plant coupled with a battery energy storage device. The results of the stochastic planner are comparable with those of the deterministic planner, even when the prediction error variance is non-negligible. However, further investigations are required to use a more realistic methodology to generate PV scenarios and a realistic controller.
- A sizing study is performed using the capacity firming framework in a different case: the ULiège case study. The results indicate that the approach is not

very sensitive to the control policy. The difference between the deterministic with perfect knowledge (or point forecasts) and stochastic with scenarios is minor. However, further investigations are required to implement a more realistic controller that uses intraday point forecasts and conduct a sensitivity analysis on the simulation parameters.

- A robust day-ahead planner is compared to its deterministic counterpart using the ULiège case study in the capacity firming framework. Overall, it allows finding a trade-off between conservative and risk-seeking policies by selecting the optimal robust optimization parameters, leading to improved economic benefits compared to the baseline.
- Finally, a stochastic approach deals with the optimal bidding of an energy retailer portfolio on the day-ahead market to assess the value of a recent class of deep learning generative models, the normalizing flows. The numerical experiments use the open-access data of the Global Energy Forecasting Competition 2014. The results demonstrate that normalizing flows can challenge Variational AutoEncoders and Generative Adversarial Networks. They are overall more accurate both in terms of quality (see Chapter 6) and value and can be used effectively by non-expert deep learning practitioners.

Chapter 15

General conclusions and perspectives

Overview

This chapter concludes the thesis and identifies four potential research directions in both forecasting and planning for microgrids. (1) Forecasting techniques of the future with the development of new machine learning models that take advantage of the underlying physical process. (2) Machine learning for optimization with models that simplify optimization planning problems by learning a sub-optimal space. (3) Modelling and simulation of energy systems by applying forecasting and decomposition techniques investigated in this thesis into open-source code. (4) Machine learning and psychology where models could identify appropriate behaviors to reduce carbon footprint. Then, inform individuals, and provide constructive opportunities by modeling individual behavior.

“*Words are loaded pistols.*”

— **Jean-Paul Sartre**

15.1 Summary

The IPCC report 'AR5 Climate Change 2014: Mitigation of Climate Change'¹, states that electrical systems are responsible for about a quarter of human-caused greenhouse gas emissions each year. According to the IPCC, the transition to a carbon-free society

¹<https://www.ipcc.ch/report/ar5/wg3/>

goes through an inevitable increase in the renewable generation's share in the energy mix and a drastic decrease in the total consumption of fossil fuels. However, a high share of renewables is challenging for power systems that have been designed and sized for dispatchable units. In addition, a major challenge is to implement these changes across all countries and contexts, as electricity systems are everywhere. Machine learning can contribute on all fronts by informing the research, deployment, and operation of electricity system technologies. High leverage contributions in power systems include [151]: accelerating the development of clean energy technologies, improving demand and clean energy forecasts, improving electricity system optimization and management, and enhancing system monitoring. This thesis focuses on two leverages: (1) the supply and demand forecast; (2) the electricity system optimization and management.

Since variable generation and electricity demand both fluctuate, they must be forecast ahead of time to inform real-time electricity scheduling and longer-term system planning. Better short-term forecasts enable system operators to reduce reliance on polluting standby plants and proactively manage increasing amounts of variable sources. Better long-term forecasts help system operators and investors to decide where and when to build variable plants. Forecasts need to become more accurate, span multiple horizons in time and space, and better quantify uncertainty to support these use cases. Therefore, probabilistic forecasts have become a vital tool to equip decision-makers, hopefully leading to better decisions in energy applications.

When balancing electricity systems, system operators use scheduling and dispatch to determine how much power every controllable generator should produce. This process is slow and complex, governed by NP-hard optimization problems [151] such as unit commitment and optimal power flow that must be coordinated across multiple time scales, from sub-second to days ahead. Scheduling becomes even more complex as electricity systems include more storage, variable generators, and flexible demand. Indeed, operators manage even more system components while simultaneously solving scheduling problems more quickly to account for real-time variations in electricity production. Thus, scheduling must improve significantly, allowing operators to rely on variable sources to manage systems.

These two challenges raise the following two central research questions:

- How to produce reliable probabilistic forecasts of renewable generation, consumption, and electricity prices?
- How to make decisions with uncertainty using probabilistic forecasts to improve scheduling?

This thesis studies the day-ahead management of a microgrid system and investigates both questions into two parts.

Part I investigates forecasting techniques and quality metrics required to produce and evaluate point and probabilistic forecasts. Then, Part II use them as input of decision-making models. Chapters 2 and 3 present the forecasting basics. They introduce the different types of forecasts to characterize the behavior of stochastic variables, such as renewable generation or electrical consumption, and the assessment procedures. An example of forecast quality evaluation is provided by employing standard deep-learning models such as recurrent neural networks to compute PV and electrical consumption point forecasts. The following Chapters of Part I study the point forecasts, quantile forecasts, scenarios, and density forecasts. Chapter 4 proposes to implement deep-learning models such as the encoder-decoder architecture to produce PV quantile forecasts. Then, a day-ahead robust optimization planner uses these forecasts in the form of prediction intervals in the second part of this thesis. Chapter 5 presents a density forecast-based approach to compute probabilistic forecasting of imbalance prices with a particular focus on the Belgian case. Finally, Chapter 6 analyzes the scenarios of renewable generation and electrical consumption of a new class of deep learning generative models, the normalizing flows. Relevant metrics assess the forecast quality. A thorough comparison is performed with the variational autoencoders and generative adversarial networks models.

Part II presents several approaches and methodologies based on optimization for decision-making under uncertainty. It employs the forecasts computed in the first part of the thesis. Chapter 8 introduces the different types of optimization strategies for decision making under uncertainty. Chapter 9 presents a value function-based approach as a way to propagate information from operational planning to real-time optimization in a deterministic framework. Chapters 10, 11, and 12 consider a grid-connected renewable generation plant coupled with a battery energy storage device in the capacity firming market. This framework has been designed to promote renewable power generation facilities in small non-interconnected grids. First, a stochastic optimization strategy is adopted for the day-ahead planner. Second, a sizing study of the system is conducted to evaluate the impact of the planning strategy: stochastic or deterministic. Finally, a robust day-ahead planner is considered to optimize the results regarding risk-aversion of the decision-maker. Finally, Chapter 13 is the extension of Chapter 6. It presents the forecast value of the deep learning generative models by considering the day-ahead market scheduling of an energy retailer. It is an easily reproducible case study designed to promote the normalizing flows in power system applications.

15.2 Future directions

We propose four primary research future directions.

15.2.1 Forecasting techniques of the future

Nowadays, the renewable energy forecasting field is highly active and dynamic. Hence, new forecasting methods will likely be proposed and used in operational problems related to power system applications in the coming years. The type of forecasts to be used as input to operational problems will depend upon the problem's nature and the formulation of the corresponding optimization problem. Machine learning algorithms of the future will need to incorporate domain-specific insights. For instance, the study [151] defines exciting research directions. We propose in the following several research directions. First, it is well known that the weather fundamentally drives both variable generation and electricity demand. Thus, machine learning algorithms should draw from climate modeling, weather forecasting innovations, and hybrid physics-plus-ML modeling techniques [171, 39]. Second, machine learning models could be designed to directly optimize system goals [45, 56]. An example is given by Donti et al. [45]. They use a deep neural network to produce demand forecasts that optimize electricity scheduling costs rather than forecast accuracy. Third, better understanding the value of improved forecasts is an interesting challenge. In this line, the benefits of specific solar forecast improvements in a region of the United States are described by Martinez-Anido et al. [120]. Finally, studying techniques that take advantage of the power system characteristics, such as graphical neural networks [177, 44]. They are capable of learning the power network structure and could provide successful contributions to hierarchical forecasting. Indeed, probabilistic graphical models reduce to Bayesian networks with a pre-defined topology and a learnable density at each node [177]. From this new perspective, the graphical normalizing flow provides a promising way to inject domain knowledge into normalizing flows while preserving Bayesian networks' interpretability and the representation capacity of normalizing flows. In this line, a graphical model is used to detect faults in rooftop solar panels [91]. This paper proposes Solar-Clique, a data-driven approach that can flag anomalies in power generation with high accuracy.

15.2.2 Machine learning for optimization

There is a broad consensus in the power system community that the uncertain nature of renewable energy sources like wind and solar is likely to induce significant changes in

the paradigms of power systems management [124]. The electricity market results from traditional practices, such as unit commitment or economic dispatch, designed given a generation mix mainly formed by dispatchable plants. Therefore, they are now to be reexamined so that stochastic producers can compete on equal terms. For instance, the capacity firming market is a new design conceived to promote renewable generation in isolated markets. The type of decision-making tools to deal with these new market designs and systems, such as microgrids, depends upon the nature of the problem itself. These tools have been studied intensively, and new methods will likely be proposed and used in power system applications. In particular, machine learning for optimization is nowadays a hot topic. It learns partially or totally the sizing space to provide a fast and efficient sizing tool. It also simplifies optimization planning problems by learning a sub-optimal space. This approach has already been investigated in a few applications. For instance:

- Machine learning can be used to approximate or simplify existing optimization problems [15, 182], and to find good starting points for optimization [92].
- Dynamic scheduling [57] and safe reinforcement learning can also be used to balance the electric grid in real-time.
- Misyris et al. [122] propose a framework for physics-informed neural networks in power system applications. In this line, Fioretto et al. [60] present a deep learning approach to the optimal power flows. The learning model exploits the information available in the similar states of the system and a dual Lagrangian method to satisfy the physical and engineering constraints present in the optimal power flows.
- Donon et al. [44] propose a neural network architecture that emulates the behavior of a physics solver that solves electricity differential equations to compute electricity flow in power grids. It uses proxies based on graph neural networks.
- Tsaousoglou et al. [167] consider an economic dispatch problem for a community of distributed energy resources, where energy management decisions are made online and under uncertainty. The economic dispatch problem is a multi-agent Markov Decision Process. The difficulties lie in the curse of dimensionality and in guaranteeing the satisfaction of constraints under uncertainty. A novel method that combines duality and deep learning tackles these challenges.

15.2.3 Modelling and simulation of energy systems

As already previously stated, the transition towards more sustainable fossil-free energy systems requires a high penetration of renewables, such as wind and solar. These new energy resources and technologies will lead to profound structural changes in energy systems, such as an increasing need for storage and drastic electrification of the heating and mobility sectors. Therefore, new flexible and open-source optimization modeling tools are required to capture the growing complexity of such future energy systems. To this end, in the past few years, several open-source models for the strategic energy planning of urban and regional energy systems have been developed. A list of such models is depicted in [115]. We select and present two² of them where we think it may be relevant to implement and test the forecasting techniques and scheduling strategies developed in this thesis.

First, EnergyScope TD [115] is a novel open-source model for the strategic energy planning of urban and regional energy systems. EnergyScope TD optimizes an entire energy system's investment and operating strategy, including electricity, heating, and mobility. Its hourly resolution, using typical days, makes the model suitable for integrating intermittent renewables, and its concise mathematical formulation and computational efficiency are appropriate for uncertainty applications. A new research direction could be integrating multi-criterion optimization to consider the carbon footprint, the land use, the energy return on investment, and the cost could be an exciting and challenging research direction [126]. Solving such a complex optimization problem may require decomposition techniques such as the ones investigated in this thesis.

Second, E4CLIM [164] is open-source Python software integrating flexibility needs from variable renewable energies in the development of regional energy mixes. It aims at evaluating and optimizing energy deployment strategies with higher shares of variable renewable energies. It also assesses the impact of new technologies and climate variability and allows conducting sensitivity studies. The E4CLIM's potential was already illustrated at the country scale with an optimal recommissioning study of the 2015 Italian PV-wind mix. A new research direction could be to adapt and develop E4CLIM at the infra-regional scale and consider the electrical grid constraints. It may imply solving complex optimization problems requiring state-of-the-art decomposition techniques.

²In both cases, machine learning investigation for optimization may be an interesting research direction to simplify and improve the models.

15.2.4 Machine learning and psychology

This last research direction goes off the beaten tracks. In my opinion, achieving sustainability goals requires as much the use of relevant technology as psychology. Therefore, one of the main challenges is not designing relevant technological tools but changing how we consume and behave in our society. In other words, technology will not save us. Nevertheless, it is the way we use technology that could help us to meet the sustainability targets. Therefore, it requires research in transdisciplinary fields to design collaborative solutions and guidances towards this goal. Psychology is a comprehensive tool to study and understand why there is a gap between rhetoric and actions. An increase in psychological research on climate change has been conducted since the last decade [34]. We select and provide three useful references among the vast literature on this topic to the interested reader. Seven categories of psychological barriers or *dragons of inaction* are studied by Gifford [68] where five strategies are proposed to help overcome them. A discussion is proposed by Nielsen et al. [129] to consider how further research can make an even more significant difference in limiting Climate Change. Finally, Abrahamse and Steg [2] analyze the effect of social influence. Social influence refers to how our behavior is affected by what other people do or think. This study compares the effectiveness of different social influence approaches to sustainability using a meta-analysis. It also presents what remains to be investigated about the mechanisms linking social influence to behavior change.

This topic raises several exciting research questions. (1) How can we consider psychology and human behavior when designing deep learning algorithms or decision-making tools? (2) How can we use algorithms to influence behavior towards sustainable goals? (3) How can we design new tools and propose relevant climate change solutions without leading to a rebound effect? (4) Do we need high-tech solutions to make a significant shift? Or is it more a matter of behavior and psychology? What is the optimal trade-off?

Some actions can meaningfully reduce each person's carbon footprint and, if widely adopted, could significantly mitigate global emissions [178]. Machine learning could help identify those behaviors, inform individuals, and provide constructive opportunities by modeling individual behavior [151].

First, by understanding how is composed our carbon footprint. Indeed, as individuals, we may lack the data and knowledge to know which decisions are most impactful regarding the carbon footprint. Therefore, machine learning can help determine an individual's carbon footprint based on their behavior and personal and household data. For instance, a natural language processing model can identify the products

we purchased from a bill or a type of transportation from email. The algorithm can also monitor household electricity and gas consumption or the food we eat. Then, making it possible to predict the associated emissions and help consumers who wish to identify which behaviors result in the highest emissions. In addition, such a machine learning model could use counterfactual reasoning to demonstrate to consumers how much their emissions would be reduced for each behavior they changed. Therefore, such information can empower people to understand how they can best help mitigate climate change through behavior change.

Second, by facilitating behavior change. Indeed, machine learning can model and cluster individuals based on their climate knowledge, preferences, demographics, and consumption characteristics [61]. Thus, it predicts who will be most responsive to new technologies and sustainable behavior change. For instance, such techniques have improved the enrollment rate of customers in an energy savings program by 2-3x [5]. It can also make the effects of climate change more tangible for consumers and help motivate those who wish to act. In this line, Konstantakopoulos et al. [107] has proposed gamification through deep learning to allow further individuals to explore their energy usage.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [2] Wokje Abrahamse and Linda Steg. Social influence approaches to encourage resource conservation: A meta-analysis. *Global environmental change*, 23(6): 1773–1785, 2013.
- [3] International Energy Agency. Net zero by 2050: A roadmap for the global energy system. Technical report, International Energy Agency, 2021.
- [4] Michael S. Albergo, Denis Boyda, Daniel C. Hackett, Gurtej Kanwar, Kyle Cranmer, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan. Introduction to normalizing flows for lattice field theory, 2021.
- [5] Adrian Albert and Mehdi Maasoumy. Predictive segmentation of energy consumers. *Applied energy*, 177:435–448, 2016.
- [6] Myles Allen, Mustafa Babiker, Yang Chen, and Heleen C. de Coninck. Ipcsr15: Summary for policymakers. Technical report, Intergovernmental Panel on Climate Change, October 2018. Drafting Authors: Myles Allen (UK), Mustafa Babiker (Sudan), Yang Chen (China), Heleen de Coninck (Netherlands/EU), Sarah Connors (UK), Renée van Diemen (Netherlands), Opha Pauline Dube (Botswana), Kristie L. Ebi (USA), Francois Engelbrecht (South Africa), Marion Ferrat (UK/France), James Ford (UK/Canada), Piers Forster (UK), Sabine Fuss (Germany), Tania Guillén Bolaños (Germany/Nicaragua), Jordan Harold (UK), Ove Hoegh-Guldberg (Australia), Jean-Charles Hourcade (France), Daniel Huppmann (Austria), Daniela Jacob (Germany), Kejun Jiang (China), Tom Gabriel Johansen (Norway), Mikiko Kainuma (Japan), Kiane de Kleijne (Netherlands/EU), Elmar Kriegler (Germany), Debora Ley (Guatemala/Mexico), Diana Liverman (USA), Natalie Mahowald (USA), Valérie Masson-Delmotte (France), J. B. Robin Matthews (UK), Richard Millar (UK), Katja Mintenbeck (Germany),

- Angela Morelli (Norway/Italy), Wilfran Moufouma-Okia (France/Congo), Luis Mundaca (Sweden/Chile), Maïke Nicolai (Germany), Chukwumerije Okereke (UK/Nigeria), Minal Pathak (India), Antony Payne (UK), Roz Pidcock (UK), Anna Pirani (Italy), Elvira Poloczanska (UK/Australia), HansOtto Pörtner (Germany), Aromar Revi (India), Keywan Riahi (Austria), Debra C. Roberts (South Africa), Joeri Rogelj (Austria/Belgium), Joyashree Roy (India), Sonia I. Seneviratne (Switzerland), Priyadarshi R. Shukla (India), James Skea (UK), Raphael Slade (UK), Drew Shindell (USA), Chandni Singh (India), William Solecki (USA), Linda Steg (Netherlands), Michael Taylor (Jamaica), Petra Tschakert (Australia/Austria), Henri Waisman (France), Rachel Warren (UK), Panmao Zhai (China), Kirsten Zickfeld (Canada).
- [7] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks, 2017.
- [8] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [9] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [10] Alberto Bemporad, Andrea Garulli, Simone Paoletti, and Antonio Vicino. A greedy approach to identification of piecewise affine models. In *International Workshop on Hybrid Systems: Computation and Control*, pages 97–112. Springer, 2003.
- [11] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton University Press, 2009.
- [12] Jacques F Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische mathematik*, 4(1):238–252, 1962.
- [13] Yoshua Bengio and Yves Grandvalet. No unbiased estimator of the variance of k-fold cross-validation. *Journal of machine learning research*, 5(Sep):1089–1105, 2004.
- [14] Dimitris Bertsimas and Shimrit Shtern. A scalable algorithm for two-stage adaptive linear optimization, 2018.
- [15] Dimitris Bertsimas and Bartolomeo Stellato. Online mixed-integer optimization in milliseconds, 2021.
- [16] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [17] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.

- [18] Dimitris Bertsimas, Eugene Litvinov, Xu Andy Sun, Jinye Zhao, and Tongxin Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE transactions on power systems*, 28(1):52–63, 2012.
- [19] Gianni Bianchini, Simone Paoletti, Antonio Vicino, Franco Corti, and Federico Nebiacolombo. Model estimation of photovoltaic power generation using partial information. In *IEEE PES ISGT Europe 2013*, pages 1–5. IEEE, 2013.
- [20] Gianni Bianchini, Daniele Pepe, and Antonio Vicino. Estimation of photovoltaic generation forecasting models using limited information. *Automatica*, 113:108688, 2020.
- [21] Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- [22] John R Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- [23] Eilyan Y Bitar, Ram Rajagopal, Pramod P Khargonekar, Kameshwar Poola, and Pravin Varaiya. Bringing wind energy to market. *IEEE Transactions on Power Systems*, 27(3):1225–1235, 2012.
- [24] Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G. Willcocks. Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models, 2021.
- [25] Trine Krogh Boomsma, Nina Juul, and Stein-Erik Fleten. Bidding in sequential electricity markets: The nordic case. *European Journal of Operational Research*, 238(3):797–809, 2014.
- [26] Ali Borji. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding*, 179:41–65, 2019.
- [27] Jérémie Bottieau, Louis Hubert, Zacharie De Grève, François Vallée, and Jean-François Toubeau. Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement. *IEEE Transactions on Power Systems*, 35(2):1218–1230, 2019.
- [28] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [29] Simon Camal, Fei Teng, Andrea Michiorri, Georges Kariniotakis, and L Badesa. Scenario generation of aggregated wind, photovoltaics and small hydro production for power systems applications. *Applied Energy*, 242:1396–1406, 2019.
- [30] Yize Chen, Pan Li, and Baosen Zhang. Bayesian renewables scenario generation via deep generative networks. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2018.
- [31] Yize Chen, Yishen Wang, Daniel Kirschen, and Baosen Zhang. Model-free renewable scenario generation using generative adversarial networks. *IEEE Transactions on Power Systems*, 33(3):3265–3275, 2018.

- [32] Zhenghua Chen and Chaoyang Jiang. Building occupancy modeling using generative adversarial network. *Energy and Buildings*, 174:372–379, 2018.
- [33] François Chollet et al. Keras. <https://keras.io>, 2015.
- [34] Susan Clayton and Christie Manning. *Psychology and climate change: Human perceptions, impacts, and responses*. Academic Press, 2018.
- [35] Seán Collins, John Paul Deane, Kris Poncelet, Evangelos Panos, Robert C Pietzcker, Erik Delarue, and Brian Pádraig Ó Gallachóir. Integrating short term variations of the power system into integrated energy system models: A methodological review. *Renewable and Sustainable Energy Reviews*, 76:839–856, 2017.
- [36] Stefano Raimondi Cominesi, Marcello Farina, Luca Giulioni, Bruno Picasso, and Riccardo Scattolini. A two-layer stochastic model predictive control scheme for microgrids. *IEEE Transactions on Control Systems Technology*, 26(1):1–13, 2017.
- [37] Abdelkader Dairi, Fouzi Harrou, Ying Sun, and Sofiane Khadraoui. Short-term forecasting of photovoltaic solar power production using variational auto-encoder driven deep learning approach. *Applied Sciences*, 10(23):8400, 2020.
- [38] Selmane Dakir and Bertrand Cornélusse. Combining optimization and simulation for microgrid sizing. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494788.
- [39] Utpal Kumar Das, Kok Soon Tey, Mehdi Seyedmahmoudian, Saad Mekhilef, Moh Yamani Idna Idris, Willem Van Deventer, Bend Horan, and Alex Stojcevski. Forecasting of photovoltaic power generation and model optimization: A review. *Renewable and Sustainable Energy Reviews*, 81:912–928, 2018.
- [40] Jan G De Gooijer and Rob J Hyndman. 25 years of time series forecasting. *International journal of forecasting*, 22(3):443–473, 2006.
- [41] Kristof De Vos. Negative wholesale electricity prices in the german, french and belgian day-ahead, intra-day and real-time markets. *The Electricity Journal*, 28(4):36–50, 2015.
- [42] Masson Delmotte, P. Zhai V., A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.). *IPCC, 2021: Summary for Policymakers*. In: *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2021.
- [43] Francis X Diebold and Robert S Mariano. Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1):134–144, 2002.
- [44] Balthazar Donon, Benjamin Donnot, Isabelle Guyon, and Antoine Marot. Graph neural solver for power systems. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019.

- [45] Priya L. Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization, 2019.
- [46] RN Dows and EJ Gough. Pvusa procurement, acceptance, and rating practices for photovoltaic power plants. Technical report, Pacific Gas and Electric Co., San Ramon, CA (United States). Dept. of Research and Development, 1995.
- [47] Laurine Duchesne. *Machine Learning of Proxies for Power Systems Reliability Management in Operation Planning*. PhD thesis, Université de Liège, Liège, Belgique, 2021.
- [48] Jonathan Dumas and Bertrand Cornélusse. Classification of load forecasting studies by forecasting problem to select load forecasting techniques and methodologies, 2020. URL <https://arxiv.org/abs/1901.05052>.
- [49] Jonathan Dumas, Ioannis Boukas, Miguel Manuel de Villena, Sébastien Mathieu, and Bertrand Cornélusse. Probabilistic forecasting of imbalance prices in the belgian context. In *2019 16th International Conference on the European Energy Market (EEM)*, pages 1–7. IEEE, 2019. URL <https://arxiv.org/abs/2106.07361>.
- [50] Jonathan Dumas, Bertrand Cornélusse, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Stochastic and deterministic formulations for capacity firming nominations. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–7. IEEE, 2020. URL <https://arxiv.org/abs/2106.02425>.
- [51] Jonathan Dumas, Colin Cointe, Xavier Fettweis, and Bertrand Cornélusse. Deep learning-based multi-output quantile forecasting of pv generation. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494976. URL <https://arxiv.org/abs/2106.01271>.
- [52] Jonathan Dumas, Colin Cointe, Antoine Wehenkel, Antonio Sutera, Xavier Fettweis, and Bertrand Cornélusse. A probabilistic forecast-driven strategy for a risk-aware participation in the capacity firming market. *IEEE Transactions on Sustainable Energy*, pages 1–1, 2021. doi: 10.1109/TSTE.2021.3117594. URL <https://arxiv.org/abs/2105.13801>.
- [53] Jonathan Dumas, Bertrand Cornélusse, Xavier Fettweis, Antonello Giannitrapani, Simone Paoletti, and Antonio Vicino. Probabilistic forecasting for sizing in the capacity firming framework. In *2021 IEEE Madrid PowerTech*, pages 1–6, 2021. doi: 10.1109/PowerTech46648.2021.9494947. URL <https://arxiv.org/abs/2106.02323>.
- [54] Jonathan Dumas, Selmane Dakir, Clément Liu, and Bertrand Cornélusse. Coordination of operational planning and real-time optimization in microgrids. *Electric Power Systems Research*, 190:106634, 2021. URL <https://arxiv.org/abs/2106.02374>.
- [55] Jonathan Dumas, Antoine Wehenkel, Damien Lanaspèze, Bertrand Cornélusse, and Antonio Sutera. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022. ISSN

- 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2021.117871>. URL <https://arxiv.org/abs/2106.09370>.
- [56] Adam N Elmachtoub and Paul Grigas. Smart “predict, then optimize”. *Management Science*, 2021.
- [57] Andreas Essl, André Ortner, Reinhard Haas, and Peter Hettegger. Machine learning analysis for a flexibility energy approach towards renewable energy integration with dynamic forecasting of electricity balancing power. In *2017 14th International Conference on the European Energy Market (EEM)*, pages 1–6. IEEE, 2017.
- [58] Tom Fawcett. Roc graphs: Notes and practical considerations for researchers. *Machine learning*, 31(1):1–38, 2004.
- [59] Xavier Fettweis, Jason Box, Cécile Agosta, Charles Amory, Christoph Kittel, Charlotte Lang, Dirk van As, Horst Machguth, and Hubert Gallée. Reconstructions of the 1900–2015 greenland ice sheet surface mass balance using the regional climate MAR model. *Cryosphere (The)*, 11:1015–1033, 2017.
- [60] Ferdinando Fioretto, Terrence W.K. Mak, and Pascal Van Hentenryck. Predicting ac optimal power flows: Combining deep learning and lagrangian dual methods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):630–637, Apr. 2020. doi: 10.1609/aaai.v34i01.5403.
- [61] Elizabeth Gabe-Thomas, Ian Walker, Bas Verplanken, and Gavin Shaddick. Householders’ mental models of domestic energy consumption: using a sort-and-cluster method to identify shared concepts of appliance similarity. *PloS one*, 11(7):e0158949, 2016.
- [62] Maria P Garcia and Daniel S Kirschen. Forecasting system imbalance volumes in competitive electricity markets. *IEEE Transactions on Power Systems*, 21(1): 240–248, 2006.
- [63] Leijiao Ge, Wenlong Liao, Shouxiang Wang, Birgitte Bak-Jensen, and Jayakrishnan Radhakrishna Pillai. Modeling daily load profiles of distribution network for scenario generation using flow-based generative network. *IEEE Access*, 8: 77587–77597, 2020.
- [64] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.
- [65] Antonio Giannitrapani, Simone Paoletti, Antonio Vicino, and Donato Zarrilli. Bidding strategies for renewable energy generation with non stationary statistics. *IFAC Proceedings Volumes*, 47(3):10784–10789, 2014.
- [66] Antonio Giannitrapani, Simone Paoletti, Antonio Vicino, and Donato Zarrilli. Bidding wind energy exploiting wind speed forecasts. *IEEE Transactions on Power Systems*, 31(4):2647–2656, 2015.

- [67] Antonio Giannitrapani, Simone Paoletti, Antonio Vicino, and Donato Zarrilli. Bidding wind energy exploiting wind speed forecasts. *IEEE Transactions on Power Systems*, 31(4):2647–2656, 2016.
- [68] Robert Gifford. The dragons of inaction: psychological barriers that limit climate change mitigation and adaptation. *American psychologist*, 66(4):290, 2011.
- [69] Tilmann Gneiting and Matthias Katzfuss. Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1:125–151, 2014.
- [70] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- [71] Tilmann Gneiting, Larissa I Stanberry, Eric P Gritmit, Leonhard Held, and Nicholas A Johnson. Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test*, 17(2):211–235, 2008.
- [72] Faranak Golestaneh, Hoay Beng Gooi, and Pierre Pinson. Generation and evaluation of space–time trajectories of photovoltaic power. *Applied Energy*, 176: 80–91, 2016.
- [73] Faranak Golestaneh, Pierre Pinson, and Hoay Beng Gooi. Very short-term nonparametric probabilistic forecasting of renewable energy generation—with application to solar energy. *IEEE Transactions on Power Systems*, 31(5):3850–3863, 2016.
- [74] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [75] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [76] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [77] Richard Green and Nicholas Vasilakos. Market behaviour with large amounts of intermittent generation. *Energy Policy*, 38(7):3211–3220, 2010.
- [78] Stephen R Green and Jonathan Gair. Complete parameter inference for gw150914 using deep learning. *Machine Learning: Science and Technology*, 2(3):03LT01, 2021.
- [79] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans, 2017.
- [80] Pierre Haessig. *Dimensionnement et gestion d’un stockage d’énergie pour l’atténuation des incertitudes de production éolienne*. PhD thesis, Cachan, Ecole normale supérieure, 2014.

- [81] Pierre Haessig, Bernard Multon, Hamid Ben Ahmed, Stéphane Lascaud, and Lionel Jamy. Aging-aware nas battery model in a stochastic wind-storage simulation framework. In *2013 IEEE Grenoble Conference*, pages 1–6. IEEE, 2013.
- [82] Pierre Haessig, Bernard Multon, Hamid Ben Ahmed, Stéphane Lascaud, and Pascal Bondon. Energy storage sizing for wind power: impact of the autocorrelation of day-ahead forecast errors. *Wind Energy*, 18(1):43–57, 2015.
- [83] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [84] Hansika Hewamalage, Christoph Bergmeir, and Kasun Bandara. Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1):388–427, 2020.
- [85] Tao Hong and Shu Fan. Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 32(3):914–938, 2016.
- [86] Tao Hong, Pierre Pinson, Shu Fan, Hamidreza Zareipour, Alberto Troccoli, and Rob J. Hyndman. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, 32(3):896–913, 2016. ISSN 0169-2070. doi: <https://doi.org/10.1016/j.ijforecast.2016.02.001>.
- [87] Tao Hong, Pierre Pinson, Yi Wang, Rafał Weron, Dazhi Yang, and Hamidreza Zareipour. Energy forecasting: A review and outlook. *IEEE Open Access Journal of Power and Energy*, 7:376–388, 2020. doi: 10.1109/OAJPE.2020.3029979.
- [88] Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron Courville. Neural autoregressive flows. In *International Conference on Machine Learning*, pages 2078–2087. PMLR, 2018.
- [89] Gao Huang, Yang Yuan, Qiantong Xu, Chuan Guo, Yu Sun, Felix Wu, and Kilian Weinberger. An empirical study on evaluation metrics of generative adversarial networks, 2018.
- [90] Semich Impram, Secil Varbak Nese, and Bülent Oral. Challenges of renewable energy penetration on power system flexibility: A survey. *Energy Strategy Reviews*, 31:100539, 2020.
- [91] Srinivasan Iyengar, Stephen Lee, Daniel Sheldon, and Prashant Shenoy. Solar-clique: Detecting anomalies in residential solar arrays. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 1–10, 2018.
- [92] Mahdi Jamei, Letif Mones, Alex Robson, Lyndon White, James Requeima, and Cozmin Ududec. Meta-optimization of optimal power flow. In *ICML Workshop, Climate Change: How Can AI Help*, 2019.
- [93] Congmei Jiang, Yongfang Mao, Yi Chai, and Mingbiao Yu. Day-ahead renewable scenario forecasts based on generative adversarial networks. *International Journal of Energy Research*, 45(5):7572–7587, 2021.

-
- [94] Quanyuan Jiang, Meidong Xue, and Guangchao Geng. Energy management of microgrid in grid-connected and stand-alone modes. *IEEE transactions on power systems*, 28(3):3380–3389, 2013.
- [95] Ruiwei Jiang, Jianhui Wang, and Yongpei Guan. Robust unit commitment with wind power and pumped storage hydro. *IEEE Transactions on Power Systems*, 27(2):800–810, 2011.
- [96] Chengquan Ju, Peng Wang, Lalit Goel, and Yan Xu. A two-layer energy management system for microgrids with hybrid energy storage considering degradation costs. *IEEE Transactions on Smart Grid*, 9(6):6047–6057, 2017.
- [97] Sami H Karaki, Bassel A Salim, and Riad B Chedid. Probabilistic model of a two-site wind energy conversion system. *IEEE Transactions on Energy Conversion*, 17(4):530–536, 2002.
- [98] SH Karaki, RB Chedid, and R Ramadan. Probabilistic performance assessment of autonomous solar-wind energy conversion systems. *IEEE Transactions on energy conversion*, 14(3):766–772, 1999.
- [99] Janina C Ketterer. The impact of wind power generation on the electricity price in germany. *Energy Economics*, 44:270–280, 2014.
- [100] Abdolrahman Khoshrou and Eric J Pauwels. Short-term scenario-based probabilistic load forecasting: A data-driven approach. *Applied Energy*, 238:1258–1268, 2019.
- [101] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [102] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014.
- [103] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [104] Gro Klæboe, Anders Lund Eriksrud, and Stein-Erik Fleten. Benchmarking time series based forecasting models for electricity balancing market prices. *Energy Systems*, 6(1):43–61, 2015.
- [105] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [106] Roger Koenker and Gilbert Bassett Jr. Regression quantiles. *Econometrica: journal of the Econometric Society*, pages 33–50, 1978.
- [107] Ioannis C Konstantakopoulos, Andrew R Barkan, Shiyong He, Tanya Veeravalli, Huihan Liu, and Costas Spanos. A deep learning and gamification approach to improving human-building interaction and energy efficiency in smart infrastructure. *Applied energy*, 237:810–821, 2019.

- [108] Ranjeet Kumar, Michael J Wenzel, Matthew J Ellis, Mohammad N ElBsat, Kirk H Drees, and Victor M Zavala. A stochastic dual dynamic programming framework for multiscale mpc. *IFAC-PapersOnLine*, 51(20):493–498, 2018.
- [109] Olivier Labussière and Alain Nadaï. *Energy transitions: a socio-technical inquiry*. Springer, 2018.
- [110] Jian Lan, Qinglai Guo, and Hongbin Sun. Demand side data generating based on conditional generative adversarial networks. *Energy Procedia*, 152:1188–1193, 2018.
- [111] Mark Landry, Thomas P Erlinger, David Patschke, and Craig Varrichio. Probabilistic gradient boosting machines for gefcom2014 wind forecasting. *International Journal of Forecasting*, 32(3):1061–1066, 2016.
- [112] Philippe Lauret, Mathieu David, and Pierre Pinson. Verification of solar irradiance probabilistic forecasts. *Solar Energy*, 194:254–271, 2019.
- [113] Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- [114] Gauthier Limpens. *Generating energy transition pathways: application to Belgium*. PhD thesis, UCL-Université Catholique de Louvain, 2021.
- [115] Gauthier Limpens, Stefano Moret, Hervé Jeanmart, and Francois Maréchal. Energyscope td: A novel open-source model for regional energy systems. *Applied Energy*, 255:113729, 2019.
- [116] Hui Lin and Halit Üster. Exact and heuristic algorithms for data-gathering cluster-based wireless sensor network design problem. *IEEE/ACM transactions on networking*, 22(3):903–916, 2013.
- [117] Haitao Liu, Jianfei Cai, and Yew-Soon Ong. Remarks on multi-output gaussian process regression. *Knowledge-Based Systems*, 144:102–121, 2018.
- [118] B Looney. Statistical review of world energy, 2020, 2020.
- [119] Gilles Louppe. *Understanding random forests: From theory to practice*. PhD thesis, Université de Liège, Liège, Belgique, 2014.
- [120] Carlo Brancucci Martinez-Anido, Benjamin Botor, Anthony R Florita, Caroline Draxl, Siyuan Lu, Hendrik F Hamann, and Bri-Mathias Hodge. The value of day-ahead solar power forecasting improvement. *Solar Energy*, 129:192–203, 2016.
- [121] Aleksei Mashlakov, Toni Karonen, Lasse Lensu, Arto Kaarna, and Samuli Honkapuro. Assessing the performance of deep learning models for multivariate probabilistic energy forecasting. *Applied Energy*, 285:116405, 2021.
- [122] George S Misyris, Andreas Venzke, and Spyros Chatzivasileiadis. Physics-informed neural networks for power systems. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.

- [123] Juan M Morales, Roberto Minguez, and Antonio J Conejo. A methodology to generate statistically dependent wind speed scenarios. *Applied Energy*, 87(3): 843–855, 2010.
- [124] Juan M Morales, Antonio J Conejo, Henrik Madsen, Pierre Pinson, and Marco Zugno. *Integrating renewables in electricity markets: operational problems*, volume 205. Springer Science & Business Media, 2013.
- [125] L Moretti, S Polimeni, L Meraldi, P Raboni, S Leva, and G Manzolini. Assessing the impact of a two-layer predictive dispatch algorithm on design and operation of off-grid hybrid microgrids. *Renewable Energy*, 143:1439–1453, 2019.
- [126] Bastien Muyltermans and Gauthier. Nève. Multi-criteria optimisation of an energy system and application to the belgian case. Master’s thesis, Catholic University of Louvain, 2021.
- [127] Arnold N’Goran. *Contrôle optimal et gestion énergétique d’une station d’énergie autonome par optimisation robuste*. PhD thesis, Université Paris sciences et lettres, 2020.
- [128] Arnold N’Goran, Bruno Daugrois, Marc Lotteau, and Sophie Demassey. Optimal engagement and operation of a grid-connected pv/battery system. In *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, pages 1–5. IEEE, 2019.
- [129] Kristian S Nielsen, Susan Clayton, Paul C Stern, Thomas Dietz, Stuart Capstick, and Lorraine Whitmarsh. How psychology can help limit climate change. *American Psychologist*, 76(1):130, 2021.
- [130] Anne Olhoff and John M. Christensen, editors. *Emissions Gap Report 2020*. UNEP DTU Partnership, 2020. URL <https://www.unep.org/emissions-gap-report-2020>.
- [131] Daniel E Olivares, Ali Mehrizi-Sani, Amir H Etemadi, Claudio A Cañizares, Reza Iravani, Mehrdad Kazerani, Amir H Hajimiragha, Oriol Gomis-Bellmunt, Maryam Saadifard, Rodrigo Palma-Behnke, et al. Trends in microgrid control. *IEEE Transactions on smart grid*, 5(4):1905–1919, 2014.
- [132] Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International conference on machine learning*, pages 3918–3926. PMLR, 2018.
- [133] Jorge Ángel González Ordiano, Lutz Gröll, Ralf Mikut, and Veit Hagenmeyer. Probabilistic energy forecasting using the nearest neighbors quantile filter and quantile regression. *International Journal of Forecasting*, 36(2):310–323, 2020.
- [134] Omid Palizban, Kimmo Kauhaniemi, and Josep M Guerrero. Microgrids in active network management part i: Hierarchical control, energy storage, virtual power plants, and market participation. *Renewable and Sustainable Energy Reviews*, 36:428–439, 2014.

- [135] George Papaefthymiou and Dorota Kurowicka. Using copulas for modeling stochastic dependence in power system uncertainty analysis. *IEEE Transactions on Power Systems*, 24(1):40–49, 2008.
- [136] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.
- [137] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- [138] Alessandra Parisio, Evangelos Rikos, and Luigi Glielmo. Stochastic model predictive control for economic/environmental operation management of microgrids: An experimental case study. *Journal of Process Control*, 43:24–37, 2016.
- [139] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [140] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [141] Fernando Pérez-Cruz. Kullback-leibler divergence estimation of continuous distributions. In *2008 IEEE international symposium on information theory*, pages 1666–1670. IEEE, 2008.
- [142] Pierre Pinson and Robin Girard. Evaluating the quality of scenarios of short-term wind power generation. *Applied Energy*, 96:12–20, 2012.
- [143] Pierre Pinson, Christophe Chevallier, and George N Kariniotakis. Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE Transactions on Power Systems*, 22(3):1148–1156, 2007.
- [144] Pierre Pinson, Henrik Aa Nielsen, Jan K Møller, Henrik Madsen, and George N Kariniotakis. Non-parametric probabilistic forecasts of wind power: required properties and evaluation. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 10(6):497–516, 2007.
- [145] Pierre Pinson, Henrik Madsen, Henrik Aa Nielsen, George Papaefthymiou, and Bernd Klöckl. From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 12(1):51–62, 2009.
- [146] S. Polimeni, L. Moretti, G. Manzolini, S. Leva, L. Meraldi, and P. Raboni. Numerical and experimental testing of predictive ems algorithms for pv-bess residential microgrid. In *2019 IEEE Milan PowerTech*, pages 1–6, 2019. doi: 10.1109/PTC.2019.8810548.

- [147] Yuchen Qi, Wei Hu, Yu Dong, Yue Fan, Ling Dong, and Ming Xiao. Optimal configuration of concentrating solar power in multienergy power systems with an improved variational autoencoder. *Applied Energy*, 274:115124, 2020.
- [148] Ragheb Rahmaniani, Teodor Gabriel Crainic, Michel Gendreau, and Walter Rei. The benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259(3):801–817, 2017.
- [149] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538. PMLR, 2015.
- [150] Yann Riffonneau, Seddik Bacha, Franck Barruel, and Stephane Ploix. Optimal power flow management for grid connected pv systems with batteries. *IEEE Transactions on sustainable energy*, 2(3):309–320, 2011.
- [151] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. Tackling climate change with machine learning, 2019.
- [152] Lars Ruthotto and Eldad Haber. An introduction to deep generative modeling. *GAMM-Mitteilungen*, page e202100008, 2021.
- [153] Julia Sachs and Oliver Sawodny. A two-stage model predictive control strategy for economic diesel-pv-battery island microgrid operation in rural areas. *IEEE Transactions on Sustainable Energy*, 7(3):903–913, 2016.
- [154] Yves-Marie Saint-Drenan. Wind power predictions analysis. part 1. tennet imbalance price system. development of a model for tennet imbalance price. Technical report, Energy research Centre of the Netherlands ECN, 2002.
- [155] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191, 2020.
- [156] Iacopo Savelli, Bertrand Cornélusse, Antonio Giannitrapani, Simone Paoletti, and Antonio Vicino. A new approach to electricity market clearing with uniform purchase price and curtailable block orders. *Applied energy*, 226:618–630, 2018.
- [157] Michael Scheuerer and Thomas M Hamill. Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Monthly Weather Review*, 143(4):1321–1334, 2015.
- [158] Heng Shi, Minghao Xu, and Ran Li. Deep learning for household load forecasting—a novel pooling deep rnn. *IEEE Transactions on Smart Grid*, 9(5):5271–5280, 2017.

- [159] Jéssica Alice A Silva, Juan Camilo López, Nataly Bañol Arias, Marcos J Rider, and Luiz CP da Silva. An optimal stochastic energy management system for resilient microgrids. *Applied Energy*, 300:117435, 2021.
- [160] Bharatkumar V Solanki, Claudio A Cañizares, and Kankar Bhattacharya. Practical energy management systems for isolated microgrids. *IEEE Transactions on Smart Grid*, 2018.
- [161] Mucun Sun, Cong Feng, and Jie Zhang. Probabilistic solar power forecasting based on weather scenario generation. *Applied Energy*, 266:114823, 2020.
- [162] Antonio Sutera. *Importance measures derived from random forests: characterisation and extension*. PhD thesis, Université de Liège, Liège, Belgique, 2021.
- [163] Souhaib Ben Taieb, Gianluca Bontempi, Amir F Atiya, and Antti Sorjamaa. A review and comparison of strategies for multi-step ahead time series forecasting based on the nn5 forecasting competition. *Expert systems with applications*, 39(8):7067–7083, 2012.
- [164] Alexis Tantet, Marc Stéfanon, Philippe Drobinski, Jordi Badosa, Silvia Concettini, Anna Cretì, Claudia D’ambrosio, Dimitri Thomopoulos, and Peter Tankov. E4clim 1.0: The energy for a climate integrated model: Description and application to italy. *Energies*, 12(22):4299, 2019.
- [165] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models, 2016.
- [166] Jean-François Toubreau, Jérémie Bottieau, François Vallée, and Zacharie De Grève. Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets. *IEEE Transactions on Power Systems*, 34(2):1203–1215, 2018.
- [167] Georgios Tsaousoglou, Katerina Mitropoulou, Konstantinos Steriotis, Nikolaos G Paterakis, Pierre Pinson, and Emmanouel Varvarigos. Managing distributed flexibility under uncertainty by combining deep learning with duality. *IEEE Transactions on Sustainable Energy*, 2021.
- [168] Falko Ueckerdt, Robert Brecha, and Gunnar Luderer. Analyzing major challenges of wind and solar variability in power systems. *Renewable energy*, 81:1–10, 2015.
- [169] Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in neural information processing systems*, pages 831–838, 1992.
- [170] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [171] Cyril Voyant, Gilles Notton, Soteris Kalogirou, Marie-Laure Nivet, Christophe Paoli, Fabrice Motte, and Alexis Fouilloy. Machine learning methods for solar radiation forecasting: A review. *Renewable Energy*, 105:569–582, 2017.
- [172] Bo Wang and Tao Chen. Gaussian process regression with multiple response variables. *Chemometrics and Intelligent Laboratory Systems*, 142:159–165, 2015.

- [173] Pu Wang, Bidong Liu, and Tao Hong. Electric load forecasting with recency effect: A big data approach. *International Journal of Forecasting*, 32(3):585–597, 2016.
- [174] Ran Wang, Ping Wang, and Gaoxi Xiao. A robust optimization approach for energy generation scheduling in microgrids. *Energy Conversion and Management*, 106:597–607, 2015.
- [175] Yi Wang, Gabriela Hug, Zijie Liu, and Ning Zhang. Modeling load forecast uncertainty using generative adversarial networks. *Electric Power Systems Research*, 189:106732, 2020.
- [176] Antoine Wehenkel and Gilles Louppe. Unconstrained monotonic neural networks. In *Advances in Neural Information Processing Systems*, pages 1545–1555, 2019.
- [177] Antoine Wehenkel and Gilles Louppe. Graphical normalizing flows. In *International Conference on Artificial Intelligence and Statistics*, pages 37–45. PMLR, 2021.
- [178] Katie Williamson, Aven Satre-Meloy, Katie Velasco, and Kevin Green. Climate change needs behavior change: Making the case for behavioral solutions to reduce global warming. *Rare: Arlington, VA, USA*, 2018.
- [179] Xiong Wu, Xiuli Wang, and Chong Qu. A hierarchical framework for generation scheduling of microgrids. *IEEE Transactions on Power Delivery*, 29(6):2448–2457, 2014.
- [180] Ran Yuan, Bo Wang, Zhixin Mao, and Junzo Watada. Multi-objective wind power scenario forecasting based on pg-gan. *Energy*, page 120379, 2021.
- [181] Michaël Zamo and Philippe Naveau. Estimation of the continuous ranked probability score with limited information and applications to ensemble weather forecasts. *Mathematical Geosciences*, 50(2):209–234, 2018.
- [182] Ahmed S Zamzam and Kyri Baker. Learning optimal solutions for extremely fast ac optimal power flow. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pages 1–6. IEEE, 2020.
- [183] Bo Zeng and Long Zhao. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461, 2013.
- [184] Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola. Dive into deep learning, 2021.
- [185] Hongxuan Zhang, Zongxiang Lu, Wei Hu, Yiting Wang, Ling Dong, and Jietan Zhang. Coordinated optimal operation of hydro–wind–solar integrated systems. *Applied Energy*, 242:883–896, 2019.

-
- [186] Yufan Zhang, Qian Ai, Fei Xiao, Ran Hao, and Tianguang Lu. Typical wind power scenario generation for multiple wind farms using conditional improved wasserstein generative adversarial network. *International Journal of Electrical Power & Energy Systems*, 114:105388, 2020.
- [187] Hongxuan Zhanga, Wei Hua, Rui Yub, Maolin Tangb, and Lijie Dingc. Optimized operation of cascade reservoirs considering complementary characteristics between wind and photovoltaic based on variational auto-encoder. In *MATEC Web of Conferences*, volume 246, page 01077. EDP Sciences, 2018.
- [188] Long Zhao and Bo Zeng. An exact algorithm for two-stage robust optimization with mixed integer recourse problems. *submitted, available on Optimization-Online.org*, 2012.
- [189] Long Zhao and Bo Zeng. Robust unit commitment problem with demand response and wind energy. In *2012 IEEE power and energy society general meeting*, pages 1–8. IEEE, 2012.
- [190] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Towards deeper understanding of variational autoencoding models, 2017.
- [191] Muhammad Fahad Zia, Elhoussin Elbouchikhi, and Mohamed Benbouzid. Microgrids energy management systems: A critical review on methods, solutions, and prospects. *Applied energy*, 222:1033–1055, 2018.
- [192] Florian Ziel and Rafał Weron. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics*, 70:396–420, 2018.