

THESIS FOR THE DEGREE OF LICENTIATE OF
ENGINEERING

**Data-driven methods for real-time
voltage stability assessment**

HANNES HAGMAR



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2020

Data-driven methods for real-time voltage stability assessment
HANNES HAGMAR

© HANNES HAGMAR, 2020.

Division of Electric Power Engineering
Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Printed by Chalmers Reproservice
Gothenburg, Sweden 2020

HANNES HAGMAR
Division of Electric Power Engineering
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Voltage instability is a phenomenon that limits the operation and the transmission capacity of a power system. An operation state close to the security limits enables a cost-effective utilization of the system but it could also make the system more vulnerable to disturbances. The transition towards a more sustainable energy system, with a growing share of renewable generation, will increase the complexity in voltage stability assessment and cause significant planning and operational challenges for transmission system operators.

The overall aim of this thesis is to develop a real-time voltage stability assessment tool which can be used to assist transmission system operators in monitoring voltage security limits and to provide early warnings of possible voltage instability. The thesis first analyzes the difference between static and dynamic voltage security margins, both theoretically and numerically. The results of the analysis show that power systems with a high share of loads with fast restoration dynamics, such as induction motors or power electronic controlled loads, may cause conventional static methods to assess the voltage security margins to become unreliable. Methods relying on a dynamic assessment of the security margin are in these circumstances more reliable.

However, dynamic assessment of voltage security margins is computationally challenging and can in most cases not be estimated in the time frame required by system operators in critical situations. To overcome this challenge, a machine learning-based method for fast and robust computing of the dynamic voltage security margin is proposed and tested in this thesis. The method, based on artificial neural networks, can provide real-time estimations of voltage security margins, which are then validated using a search algorithm and actual time-domain simulations. The two-step approach is proposed to mitigate any inconsistency issues associated with neural networks under new or unseen operating conditions.

Finally, a new method for voltage instability prediction is developed. The method is proposed to be used as an online tool for system operators to *predict* the system's near-future stability condition given the current operating state. The method uses a more advanced neural network based on long-short term memory. The results from case studies using the Nordic 32 test system show good performance and the network can accurately, within only a few seconds, predict voltage instability events in almost all test cases.

Keywords: Dynamic voltage security margin, voltage stability, security assessment, machine learning, neural networks, voltage instability prediction.

Acknowledgements

I would like to start off by thanking my supervisor Anh Tuan Le and my examiner Ola Carlson for very good discussions and guidance throughout this period. A special thanks go to my co-supervisor Robert Eriksson at Svenska kraftnät who have been very supportive throughout the research work, both by providing real models for me to work with, and whose experience and guidance have been most helpful. Another thanks goes to Mattias Persson at RISE for his assistance both as a co-supervisor and for helping out with the dreadful user-written models in PSS®E. I would also like to thank all the members in the reference group that have guided me and the project during these years. The financial support provided by Energimyndigheten and Svenska kraftnät is also gratefully acknowledged.

I would like to give a shout-out to my former colleagues at RISE. Especially to Anders Lindskog, who is not only a very good colleague but also the one that gave me my first real job and got me interested in research in power systems. A big thanks also to all my colleagues at the Division of Electric Power Engineering. A special recognition goes to Ankur Aloo Srivastava and Ali Fotouhi, my previous comrades-in-arms in our basement-office at Chalmers. Now, after advancement to the second-floor offices, I would also like to thank my new (and always cheerful) roommates Ioannis Bouloumpasis and Kyriaki Antoniadou-Plytaria.

Finally, my family. I planned to not be too much of a cliché here, but you are the best. Love you all. And Lina, thank you for being the best girlfriend imaginable.

Hannes Hagmar, Gothenburg, April 2020

List of Acronyms

ATC	Available Transfer Capacity
CPF	Continuation Power Flow
DSA	Dynamic Security Assessment
DT	Decision Tree
DVSM	Dynamic Voltage Security Margin
HVDC	High Voltage Direct Current
JAD	" <i>Just After Disturbance</i> "
LIVES	Local Identification of Voltage Emergency Situations
LSTM	Long Short-term Memory
LTC	Load Tap Changer
ML	Machine Learning
MSE	Mean Squared Error
NN	Neural Network
NTC	Net Transfer Capacity
OC	Operating Condition
OEL	Over-Excitation Limiter
PCLL	Post-contingency Loadability Limit
PMU	Phasor Measurement Unit
PV	Photovoltaic
QSS	Quasi-steady state
ReLU	Rectified Linear Unit Function
RNN	Recurrent Neural Network
RVSAT	Real-time Voltage stability assessment
SCADA	Supervisory Control And Data Acquisition
SOL	Secure Operating Limit
SPS	System Protection Scheme
Svk	Svenska Kraftnät
Tanh	Hyperbolic Tangent Function
TE	Thévenin Equivalent
TDT	Temporal Decision Tree
TRM	Transmission Reliability Margin
TTC	Total Transfer Capacity
VID	Voltage Instability Detection
VIP	Voltage Instability Prediction
VSA	Voltage Stability Assessment
VSI	Voltage Stability Indicators
VSM	Voltage Security Margin
WAMS	Wide Area Monitoring System



Contents

1	Introduction	1
1.1	Background	1
1.2	Problem overview	2
1.3	Previous work	5
1.4	Aim of the thesis	6
1.5	Main contributions	6
1.6	List of publications	7
1.7	Thesis outline	8
2	Voltage stability and security assessment methods	9
2.1	Definition and classification	9
2.2	Voltage stability assessment	10
2.2.1	Security and security margin assessment	12
2.2.2	Voltage instability detection (VID)	14
2.3	An overview of VID methods	16
2.3.1	Thévenin equivalent matching	17
2.3.2	Line VID methods	18
2.3.3	LIVES and new LIVES	18
2.3.4	Thévenin equivalent coupled single-port methods	19
2.3.5	Reactive reserves monitoring	19
2.3.6	Sensitivities with OEL anticipation	20
2.3.7	Machine learning-based VIP	20
3	Neural networks and machine learning	21
3.1	Machine learning: an overview	21
3.2	Feed-forward neural networks	22
3.2.1	The forward pass	23
3.2.2	Loss and learning	23
3.2.3	Overfitting and validation	24
3.2.4	Hyperparameters and network depth	25
3.3	Recurrent neural networks	25
3.3.1	Long short-term memory	26
4	Description of a real-time voltage stability assessment tool	29
4.1	Real-time voltage stability assessment tool	29
4.2	Test system	31

5	Comparison of dynamic and static voltage security margins	35
5.1	Theoretical comparison of DVSM and VSM	35
5.1.1	Small test system	35
5.1.2	Estimating DVSM and VSM	36
5.1.3	Load response after a disturbance	37
5.1.4	Transient P-V curves and fast load dynamics	38
5.1.5	DVSM versus VSM	40
5.2	Numerical comparison of DVSM and VSM	41
5.2.1	Methodology for numerical comparison	41
5.2.2	Results and discussion	43
6	Fast Dynamic Voltage Security Margin Estimation	47
6.1	Introduction	47
6.2	Methodology for fast estimation of the DVSM	47
6.2.1	Generation of training data	48
6.2.2	Architecture of the neural networks	50
6.2.3	Training	51
6.2.4	Fast DVSM estimation and dual binary search	52
6.3	Results and discussions	54
6.3.1	Regression and classification accuracy	54
6.3.2	Computational efficiency	56
6.3.3	Impact of sudden topology change	57
6.3.4	Discussions and practical applications	58
7	Voltage instability prediction using a long-short term memory network	61
7.1	Introduction	61
7.2	Methodology for LSTM-VIP	62
7.2.1	Generation of training data	62
7.2.2	Architecture of the LSTM network	66
7.2.3	Training the LSTM network	67
7.2.4	Interpretation and intuition of the VIP problem	69
7.3	Results	69
7.4	Sensitivity analysis	71
7.4.1	Impact of sequence length	71
7.4.2	Impact of measurement update rate	74
7.4.3	Generalization capability and training set requirement	75
7.5	Practical applications and requirements	76
8	Conclusions and future research	79
8.1	Conclusions	79
8.2	Future research	81
	Bibliography	83
A	Dynamic parameters for the IEESGO model	I

CHAPTER 1

Introduction

This chapter provides a background and a problem overview of the thesis, along with an overview of previous work. Furthermore, the aim and the contributions are presented together with a summary of the publications on which the thesis is based upon

1.1 Background

The transition towards a more sustainable energy system, with a growing share of renewable generation, increases the complexity in operating and controlling an electric power system. Globally, both wind and solar power are growing fast, a trend that is likely to continue given that renewable energy sources are now becoming the lowest-cost source of new power generation in most parts of the world [1, 2]. Sweden has seen a similar development, where the installed wind power capacity increased from 241 MW to 8 984 MW between the years of 2000 and 2019 [3, 4]. Simultaneously, electric power generated from conventional fossil-fueled plants needs to be decommissioned at a rapid rate to reduce carbon emissions to the atmosphere. Sweden lacks significant generation from fossil-fueled plants, and have instead relied on nuclear power for a large share of its base power. However, due to the lack of economic profitability, nuclear power is now gradually being decommissioned. By the end of 2019, reactor 2 of the Swedish nuclear power station Ringhals was shut down, and in 2020 a second reactor is planned to be decommissioned [5]. Out of originally twelve reactors, then only six will remain in operation in 2021, with a total rated capacity of about 6 900 MW.

Although a larger share of renewable energy sources in the power system is desirable from a sustainability perspective, the intermittent nature of these may cause significant planning and operational challenges. In particular, maintaining what is known as the power balance - that the generated power matches the load demand at every time instant - will become increasingly challenging in a future with a larger share of renewable generation [5]. In Sweden, the power balance is generally main-

tained by regulating the output of hydropower stations. In cases of negative power balance, and when the balancing capacity of the hydropower is not sufficient, electric power is often imported from neighboring countries. Market integration and electricity connections linking Sweden to Europe also allows ancillary services to be supported when required. In Europe, a recent example is the development of the joint automatic frequency restoration reserve, known as the European Platform for the International Coordination of the Automatic frequency restoration process and Stable System Operation [6].

To ensure that the balancing capacity of hydropower stations can be utilized to its full extent, or that sufficient power that can be imported from neighboring countries, the importance of sufficient *transmission capacity* in the power system is expected increase in the future [7]. An increased transmission capacity also allows energy to be generated where it is most economically efficient to do so and sets the limit on how much power can be transmitted through the transmission lines. Electric power systems are generally operated according to the $N-1$ contingency criterion, meaning that the system should be able to withstand the loss of any single system component, such as generation or transmission capacity, and still remain stable. A system that satisfies this criterion is said to be *secure*, and a system's transfer capacity is computed to always ensure that this criterion is fulfilled [5].

There are several stability-related phenomena that set the limit for the operation of the power systems [8]. This thesis deals with voltage stability, which refers to the capability of a system to maintain system voltages following a disturbance. Voltage stability is especially an issue in power systems where generation and loads are geographically separated over large distances; Sweden being a typical example. A voltage collapse is referring to a sequence of events accompanying voltage instability that leads to abnormally low voltages or a blackout in a significant part of a power system [8]. Although voltage collapses occur relatively seldom, they are related to extremely high costs to society and system operators need to continuously operate their systems to minimize the risk of such events.

1.2 Problem overview

While an operation close to the security limits enables a cost-effective utilization of the system, it may also make it more vulnerable to disturbances. Consequently, there exists a balance between a system that is operated efficiently and one that is operated securely. To always ensure that the $N-1$ contingency criterion is fulfilled, system operators rely on continuous estimations of the security margin for the system [9]. An accurate and fast estimation of these security margins will allow system operators to operate their systems with the confidence that the system can withstand credible disturbances.

In the case of more severe events, such as the occurrence of simultaneous multiple disturbances, system operators have to rely on system protection schemes (SPS) to

control the system back to a stable state [10]. Following such severe events, system operators need to identify whether the current system operating condition is stable, or if it is drifting towards instability. The SPS can act on signals given by voltage instability detection (VID) methods, where the onset of instability, rather than its consequences, is aimed to be detected [10]. Due to equipment in electric power systems, such as load tap changing transformers (LTCs), overexcitation limiters (OELs), and other load restoration dynamics, the time frame of a typical voltage collapse can range from a couple of seconds up to even several minutes. In the case of a possible emergency event, the ability for system operators to perform VID and act quickly and with the correct control measures is thus imperative.

To provide system operators with real-time estimates of both the security margins and to provide detection of instability, static assessment methods are generally used. The voltage security margin (VSM) is estimated by system operators to evaluate the loadability margin of a post-contingency configuration of the system [11]. The VSM is generally estimated by static assessments of the system where the stability limit up to the system collapse point is estimated [12,13]. For VID, monitoring the system voltages is the most simple method available to system operators. However, due to the reactive power support of generators and other voltage control devices in the system, the system voltages can remain at near-nominal levels a relatively long time following a disturbance. Fast system degradation and drops in system voltages may then be triggered by, for instance, the activation of field current limiters by OELs. The remaining time before a development of a potential voltage collapse may then be too short for system operators to initiate the required SPS. Other static methods for VID suffer the same dilemma; that the system degradation needs to be relatively well-developed for the VID methods to actually be able to detect instability.

Electric power systems are also becoming increasingly complex, with numerous dynamic components such as nonlinear loads, converter-based generators, and other power electronic devices found in, for instance, HVDC and flexible AC transmission systems [14,15]. These components have in common that their dynamic response following a disturbance in the system is significantly faster when compared to more conventional equipment. Depending on tuning and specified grid codes, the faster dynamic response is often beneficial for the system stability, where, for instance, the fast reactive power response of converter-based wind turbines or HVDC can help to stabilize the system following a disturbance. However, other components, such as power electronic loads, may inhibit characteristics that may cause instability events to develop faster, causing existing monitoring and emergency systems to act too slowly. The dynamic response following a contingency cannot be captured using static assessments, and studies have indicated that these might be insufficient when estimating the actual security margin [11,16]. Faster dynamics will also result in faster voltage instability events, increasing the need for faster methods for voltage instability detection.

Most power systems are monitored through supervisory control and data acquisition (SCADA) systems, in which loosely synchronized scalar measurements with a refresh-rate of around two to four seconds taken at remote terminal units are

processed by a state estimator to develop the most probable state of the power system [17]. Phasor measurement units (PMUs) are highly accurate measuring devices, allowing time-synchronized real-time phasor measurements of electrical quantities in the power system at a much higher refresh-rate than that of SCADA. If the system is fully observable by time-synchronized phasor measurements, the measurements can be filtered through a linear state estimator, allowing significantly more frequent estimates of the system state [18]. With sufficient deployment, PMUs could help in the decision making process in real-time power system operation and control by visualizing the system condition and stability margins in real-time.

The complexity of the power system requires system operators to often rely on simplified static models of how the system behaves during various operating conditions and with respect to disturbances. There are generally more advanced models available, but even with recent progress in high-performance computing, these methods are not readily available for use in real-time monitoring of a large power system. To overcome this issue, various data-driven methods and machine learning (ML) methods have been proposed in the literature. The main advantage of ML methods is that high-cost computations can be performed in an off-line setting. Once a ML algorithm is trained, it can almost instantaneously provide estimations and warnings to operators that otherwise would require time-consuming computations. However, despite years of research, examples where ML methods have been practically applied in system operators' monitoring and control systems are, to the author's best knowledge, very few. These methods still suffer from robustness issues and uncertainty when handling operating conditions not included in the training of the algorithms. Blackouts and other major failures are related to extremely high costs, and from a system operator's view, an inferior method that always works is generally preferred to a superior method that in some instances does not.

Based on the problem overview and the motivation of the thesis, the following research questions have been identified and will be examined within this thesis:

- **Research question 1:** How will the transition towards a higher penetration converter-based renewable generation and load in the power system affect available methods for security margin assessment?
- **Research question 2:** How can the current methods for assessment of security margins be improved, taking into account assessment speed, estimation accuracy, and robustness?
- **Research question 3:** What are the requirements for future voltage instability detection methods? How can they be developed to be both sufficiently advanced to capture the intricate dynamics during a voltage collapse, while at the same time be fast enough to be used in real-time?
- **Research question 4:** In the case of data-driven methods, what are the main practical aspects to consider to mitigate inconsistency problems and to increase the robustness of the methods?

In the following section, an overview of previous work related to the topics and the developed research questions in this thesis is presented.

1.3 Previous work

To ensure that a power system can handle the dynamic event following a disturbance, system operators often use an approach called dynamic security assessment (DSA). In DSA, time-domain analysis is often used to test the power system's dynamic response after a set of contingencies to ensure its ability to reach a stable post-disturbance operating point [9]. Assessment of dynamic stability is a complex task and to overcome this issue, various machine learning (ML) methods have been proposed. Examples of DSA methods based on ML are found in [15, 19–23], where mainly various decision tree (DT) or neural network (NN) methods are utilized.

In DSA, system operators are only provided with information regarding whether the *current* operating point is dynamically secure. An alternative measure of the *margin* to the most stressed point where the system can remain dynamically secure following a disturbance is the dynamic voltage security margin (DVSM). The DVSM, also referred to as the secure operating limit, is the margin to the most stressed pre-contingency operating point that can withstand a set of credible contingencies [11, 14]. Due to practical difficulties in estimating the DVSM, it has received comparatively little interest in the literature. DVSM estimation requires several time-domain simulations to trace the security limit for a set of different contingencies, which is not feasible to perform in the time frame needed by system operators. In [11], an attempt to reduce the computational cost in estimating the DVSM was developed based on using quasi-steady-state (QSS) simulations. The method was further developed in [24], where a combination of QSS and time-domain simulations was proposed to include the impact of short-term effects during the transient state following a disturbance. Although this approach reduces the computational effort compared to a full time-domain simulation, it may still prove too slow for some real-time applications. In [25–28], different ML approaches based on NNs were proposed to allow real-time estimation of the DVSM.

Following more severe events and disturbances, system operators need to identify whether the current system state is stable, or if it is drifting towards instability. There has been significant research involving different VID methods, where the complexity and the scope of the developed methods vary significantly [10]. The drawback of many developed VID methods is that they provide a late indication of instability, which reduces the available time system operators have to control the system back into stable operation and increases the risk of a voltage collapse. An alternative approach is voltage instability prediction (VIP), where the future state of the system is *predicted* using information and measurements gathered in the (short) time left following a voltage instability phenomenon.

A method for VIP based on ML was first proposed in [19], where a DT was trained

on a generated database consisting of the intermediate, short-term equilibrium that follows a disturbance. This post-contingency state, where the majority of the electromechanical transients have died out, was referred to as the "just after disturbance" (JAD) state. Extensions of the method utilizing phasor measurements have later been proposed in [29–31], where the performance of different attributes or input data have been tested. A method based on training a NN to online monitor voltage security was proposed in [32]. An attempt to incorporate some time-related features to improve the performance in VIP was presented in [33], where a temporal decision tree (TDT) approach was proposed, where some time-related features were included to improve the performance in VIP. The TDT method, further discussed in [34] and [35], could incorporate some time-related variables, such as the difference between two measurements for a specific value of elapsed time (Δt).

1.4 Aim of the thesis

The aim of this thesis is to develop a new real-time voltage stability assessment tool (RVSAT) that can support system operators and allow more efficient utilization of the transmission grid. This is achieved partly by developing methods that can provide better knowledge of the *actual* security margins in the systems, which would allow the transmission reliability margins in the system to be reduced. The thesis also aims at developing better methods to detect and assess system conditions and disturbances that might cause the system to become unstable. Fast detection of voltage instability reduces the risks and the related costs of controlling the system back into stable operation, and also allows system operators to operate their systems with higher confidence that instability can be detected quickly.

1.5 Main contributions

The main contributions of this thesis are the following.

1. A survey of current methods in VSM estimation and for VID is performed. Furthermore, the principles of the DVSM and the differences to the static VSM are illustrated using a concept called transient P - V curves to allow better interpretation of the methods. The circumstances when a dynamic security margin is preferred to a static margin are established and discussed.
2. A method for fast and robust computing of the DVSM is proposed and tested. The method is based on using ML to support the estimation of the DVSM, which otherwise is too time-consuming to perform in real-time. To mitigate inconsistency issues associated with ML methods under new or unseen operating conditions, a method to quickly validate the estimated results is developed.
3. A method for VIP is developed using a recurrent neural network with long short-term memory. This specific design of the network can utilize previous

measurements and information, such as the trend of bus voltage magnitudes, tap changes, or fault locations, to improve the accuracy for VIP. The performance of the method is compared to other structures of neural networks where the performance of the sequence-based method is evaluated.

4. A new training approach of the VIP algorithm is developed to provide system operators with an *online* assessment tool. As time progresses after a voltage instability event, the network is capable of incorporating new observations and continuously updating the assessment. The method will also allow system operators to pinpoint *where* the weakest areas in the system are located, allowing local and more cost-effective control measures.
5. A methodology for including consecutive contingencies ($N-1-1$) into the training data for the VIP algorithm is presented. Furthermore, the ability of the developed VIP method to also generalize under $N-1-1$ contingencies is examined. Such ability is especially valuable in overcoming the combinatorial increase of complexity in training.

1.6 List of publications

Following is a list of publications included in the thesis.

Paper I H. Hagmar, R. Eriksson, L. A. Tuan, "Fast Dynamic Voltage Security Margin Estimation: Concepts and Development," accepted to *IET Smart Grids Special Issue: Machine Learning in Power Systems*, 2020.

Paper II H. Hagmar, L. Tong, R. Eriksson, L. A. Tuan, "Voltage Instability Prediction Using a Deep Recurrent Neural Network," submitted to *Transactions on Power Systems*, 2020.

Paper III H. Hagmar, L. A. Tuan, O. Carlson, R. Eriksson, "On-line Voltage Instability Prediction using an Artificial Neural Network," in *Proc. 2019 PowerTech Milan*, Milan, Italy, 2019.

Paper IV H. Hagmar, L. A. Tuan, O. Carlson, R. Eriksson, "A Survey of Voltage Stability Indicators Based on Local Synchronized Phasor Measurements," in *Proc. 2018 North American Power Symposium (NAPS)*, Fargo, ND, 2018.

Additional paper produced in project but not included in the thesis:

H. Hagmar, L. A. Tuan, O. Carlson, R. Eriksson, "Integration Aspects of Full Converter Wind Turbines and the Impact on Long-term Voltage Stability," in *Proc. 2019 IEEE Power and Energy Society General Meeting*, Atlanta, GA, USA, 2019.

1.7 Thesis outline

The thesis is organized as follows:

Chapter 2 develops the definitions used in the rest of the thesis and also provides an overview of current practices in voltage security and stability assessment.

Chapter 3 provides a theoretical background on the concepts of ML and neural networks. The concepts of recurrent neural networks with long-short term memory is introduced.

Chapter 4 introduces the overall functionality of the proposed real-time voltage stability assessment tool. Furthermore, the test system used in evaluating the developed RVSAT is presented and its general characteristics are briefly discussed.

Chapter 5 presents both a theoretical and a numerical comparisons of the static and dynamic voltage security margins under various load configurations.

Chapter 6 introduces the concept of dynamic voltage security margin and discusses and illustrates the difference to conventional VSM using transient P - V curves. A method for fast estimation of the DVSM is then proposed and tested.

Chapter 7 presents a voltage instability prediction method based on a recurrent neural network using long-short term memory. The method allows system operators to continuously assess and predict whether the present system state is stable or will evolve into an alert or an emergency state in the near future.

Chapter 8 highlights the key conclusions of the thesis and provides ideas for future research.

All figures included in this thesis have been created by the author.

CHAPTER 2

Voltage stability and security assessment methods

This chapter describes the definitions and criteria with respect to voltage stability that are used in this thesis. Furthermore, the general drivers for voltage instability are discussed along with a presentation of the current practices in voltage stability and security assessment. A brief overview of the different methods used in voltage instability detection is presented. Parts of the chapter are based on the summary and results established in Paper IV.

2.1 Definition and classification

Power system stability is generally classified by the most common system variables in which instability can be observed, namely; voltage stability, rotor angle stability, and frequency stability. According to a definition by the IEEE and CIGRE Joint Task Force, voltage stability specifically refers to:

"...the ability of a power system to maintain steady voltages at all buses in the system after being subjected to a disturbance from a given initial operating condition. It depends on the ability to maintain/restore equilibrium between load demand and load supply from the power system." [36].

Voltage instability in a power system may lead to loss of loads or disconnection of other components such as tripping of transmission lines or generators due to too low voltage levels. The concept of *voltage collapse* is generally referred to as a sudden event, often initiated by a larger disturbance or by a sequence of events, leading to blackout or abnormally low voltages in the whole, or larger parts, of a power system. Historic events of voltage collapses are relatively few, but the related costs to society are extremely high [9].

The voltage stability phenomenon is complex and often requires a full network representation for its analysis [36]. Reactive power plays an important role in voltage stability, and voltage stability problems are often related to the incapability of the system to provide sufficient reactive power to keep system voltages at nominal levels. However, the *main* driver for voltage instability in power systems are the loads and the load restoration that follows a disturbance [9]. Historically, the majority of all voltage instability incidents experienced so far have resulted from larger disturbances, such as the loss of generation or transmission capacity. Following such an event, generation is redirected and the remaining transmission lines have to carry a larger current, causing increased active and reactive power losses and voltage drops. The reduced system voltages initially affect voltage-dependent loads, which are then restored by the action of, for example, LTCs, distribution voltage regulators, action of motor slip adjustments, or thermostatic load restoration. The restored loads will further increase the active and reactive power flow through a mainly inductive transmission system, which increases the reactive power losses and further deteriorates the system voltages. The system may eventually collapse when the load dynamics attempt to restore the loads beyond the capability of the generators and the transmission network.

The above-mentioned example is generally referred to as a *large-disturbance voltage stability event*. Less common are voltage instability caused by small-disturbances, such as incremental changes in system load. Since system operators continuously ensure that sufficient margins are kept to fulfil the $N-1$ contingency criterion, small-disturbance events are very seldom the sole cause of voltage instability. However, the methodology used in the small-disturbance analysis is often valuable in the analysis of the voltage security margins, which is further discussed in Section 5.1.4.

2.2 Voltage stability assessment

Voltage stability assessment (VSA) refers to the daily operational and planning activities of system operators to ensure a secure and stable operation of the power system. The main goal is to continuously ensure a *secure normal (operating) state* in the system. A system is said to be operated in a normal state if both the *load constraints* and the *operating constraints* are satisfied [9]. The load constraints ensure that the load demand is met by the generation in the system, while the operating constraints ensure that minimum or maximum limits in terms of variables such as line currents and bus voltages are satisfied.

In the event of a disturbance in the system, the system may either settle to (i) a (new) secure operating state, (ii) an (unsecure) alert state, (iii) or end up in an emergency state. Ensuring that a power system can be operated securely with respect to *all* possible disturbances would not be feasible. In practice, the power system security is assessed with respect to a set of credible disturbances, generally referred to as contingencies. Power systems are almost exclusively operated according to the $N-1$ contingency criterion, meaning that the system should be able to withstand the loss

of any single component, such as transmission or generation capacity, without the system entering an emergency state. A system capable to handle such an event without entering an emergency state is said to be *secure*. A system that cannot handle such an event without entering an emergency state is said to be operated in an *unsecure* state.

With reference to dynamical system theory, an operating state (or equilibrium) x^* is called *stable* if all solutions with an initial condition close to x^* remain near x^* for all time [9]. An equilibrium that is *not* stable is called *unstable*. Thus, a system can be stable at its current operating state, while at the same time not fulfilling the security criterion of handling a larger disturbance and still maintaining stability. An unstable system can be exemplified by a power system drifting towards voltage instability, driven by the action of load restoration. It should be noted that security and stability are sometimes used interchangeably in the literature. For instance, large-disturbance stability is in some settings defined as the same as a security. For the remaining part of the thesis, the above former stated definitions will be adopted.

The different operating states are illustrated in Fig. 2.1 and in Fig. 2.2, where the state space of a power system is reduced and illustrated in only two dimensions. In Fig. 2.1, the system is operated in the secure region. A disturbance causes the system state to transition to the unsecure region. Here, although possibly not meeting all the load constraints, the system is still stable. The system can then be transferred back to the secure region through restorative control actions and sufficient remedial actions. Alternatively, in Fig. 2.2, the system is operated in an unsecure state. Following a disturbance, the system becomes unstable and enters an emergency state, and without sufficiently fast counteractions, a larger system collapse may be imminent.




Figure 2.1: An example of a normal operating state, followed by a disturbance. Through restorative control, the system operating state is restored to normal.




Figure 2.2: An example of an unsecure operating state, followed by a disturbance. The system state ends up in an emergency state, and without sufficient emergency control actions, the system will collapse.

Generally, VSA can be divided into two different but complementary lines of defense used to avoid voltage instability: *preventive* and *emergency*. The preventive methods are mainly used during normal operation to ensure that the power system is operated securely according to the $N-1$ contingency criterion. In the event of more severe incidents or multiple disturbances, the emergency applications of VSA are required. The aim of these methods is to perform VID, allowing system operators to identify imminent voltage instability and trigger fast remedial actions. In the following sections, the main practices in these two applications of VSA are presented.

2.2.1 Security and security margin assessment

Security assessment involves assessing a power system's ability to undergo disturbances for a given operating point. The level of accuracy in the assessments vary in detail, where simpler methods are based on static assessment methods and assessing the post-contingency long-term equilibrium [9]. A static load flow assessment of a power system will have no solution if there is an absence of a post-contingency long-term equilibrium, which provides a simple way to check whether an equilibrium of the system exists. However, this simple approach suffers from several drawbacks. An absence of a post-contingency long-term equilibrium may also be the result of purely numerical problems that are not related to voltage instability. Other drawbacks include the lack of information regarding the nature and the location of the problem [9].

To overcome these issues, and to also ensure that a power system can handle the dynamic event following a large disturbance, system operators often use an approach

called dynamic security assessment (DSA). DSA refers to the analysis to determine whether or not a power system can meet security and reliability criteria in both transient and steady-state frames [37]. Commonly, time-domain analysis is used to test the power system's dynamic response for a range of different contingencies [15]. In general, DSA includes assessing not only voltage stability, but also any criteria such as thermal overloading, transient stability, or frequency stability.

More than just ensuring that the *current* operating condition is secure, the preventive assessments should also provide information to system operators with the *margin* to instability. Security margins are related to system stress, generally in terms of an increased transfer of active power. The security margins are also characterized by the direction of system stress in the parameter space by load increases and generation scheduling of various buses [9]. The voltage security margin (VSM), also referred to as the post-contingency loadability limit (PCLL), can be evaluated by estimating the loadability of a post-contingency operating point. The PCLL is commonly estimated by static assessments of the system, where parameterized continuation methods, referred to as continuation power flow (CPF), are commonly used to trace the stability limit up to the system collapse point [12, 13].

An alternative approach to estimate the security margin is by computing what is generally denoted as the secure operating limit (SOL). The SOL provides the margin to the most stressed *pre-contingency* operating point that can still withstand a set of credible contingencies [11, 14]. However, computing the transition from a *pre-contingency* operating point to a *post-contingency* operating point is numerically difficult using static assessments when the system is close to the system collapse point. Thus, the SOL requires either time-domain simulations or QSS estimations to simulate the dynamic response following a disturbance.

The SOL provides a security margin with respect to not only voltage instability, but also other stability related phenomena such as rotor angle stability or inter-area oscillations. In the remaining part of this thesis, which specifically concerns the assessment of voltage stability, the more specific term of *dynamic voltage security margin* (DVSM) will be adopted. An additional distinction between the SOL and the DVSM is that in this thesis, the DVSM is always assumed to be estimated using full time-domain simulations, whereas the SOL can also be computed using QSS simulations, or by combinations of both approaches.

The security margins are used to compute the transfer capacities in a power system, which sets the limit for how much power can be transmitted through the system. The different margins and capacities are illustrated in Fig. 2.3 with the following definitions of the notations [38]:

- **Total transfer capacity (TTC)** is the maximum transmission of active power between different areas/subsystems which is permitted with respect to a given security criterion (most commonly, the $N-1$ criterion).
- **Net transfer capacity (NTC)** is defined as: $NTC = TTC - TRM$ and is the maximum exchange possible between two areas when taking into account




Figure 2.3: Different measures of transfer capacities and reliability margins under changing operating conditions.

the uncertainties in the TTC estimation.

- **Transmission reliability margin (TRM)** is the security margin that takes into account the uncertainties on the computed TTC values. The TRM is in some cases arbitrary determined, but is generally based on [38]:
 1. Unintended deviations of physical flows during operation due to the physical functioning of load-frequency regulation.
 2. Emergency exchanges between system operators to cope with unexpected unbalanced situations in real-time.
 3. Inaccuracies, e. g. in data collection, models, and measurements.
- **Available transmission capacity (ATC)** is the maximum incremental transfer capacity possible between two parts of a power system without violating the security margins.

The ability to accurately assess and compute the transfer capacities and the related security margins is of very high importance to a system operator. Different methods for estimating the security margins in a power system, as well as comparisons between static and dynamic estimation approaches, are presented in Chapter 5.

2.2.2 Voltage instability detection (VID)

The previously described methods fall under the category of *preventive* VSA methods. In the case of more severe incidents, such as the occurrence multiple simultaneous contingencies, or in the event that the preventive methods have not been

accurate or fast enough, the system operators have to rely on emergency methods to quickly detect instability, to allow time to control the system back into a normal operation again. The aim of VID methods is thus to detect the onset of instability itself, rather than its consequences [10].

Probably the most simple, yet still commonly used, method for VID is to monitor voltage magnitudes in the power system [10]. If the voltages drop below a certain threshold value, it can provide an indicator that the system may be headed towards a collapse. However, using system voltages as an indicator of voltage instability can be problematic. In Fig. 2.4, the evolution of two transmission bus voltages are presented for the Nordic32 test system, following a disturbance. The Nordic32 test system is further presented in Section 4.2. The disturbance of concern is a tripped transmission line between buses 4032 and 4044. The voltage instability is gradually developed after the disturbance, driven by components such as LTCs and OELs. The voltage magnitudes are relatively high following the disturbance, but at some point, the mechanism of load power restoration causes the system to deteriorate to such a degree that the total power consumed in the system is reduced instead of restored for each attempt of load restoration. In Fig. 2.4, the problem of using voltage magnitudes as indicators for voltage instability is illustrated. Through the reactive power support of generators and other voltage control devices, the voltage magnitudes are kept at a relatively high level for a long time after the actual disturbance.

To overcome this issue, there has been research in the development of new and faster methods for VID. The development in phasor measurement technology also allows several new perspectives and methods for wide-area monitoring and control of power systems. The methods differ in their approach, both in complexity and the required level of system observability. An overview of common methods used in VID is presented in Section 2.3.




Figure 2.4: Evolution of bus voltages during a voltage instability event.

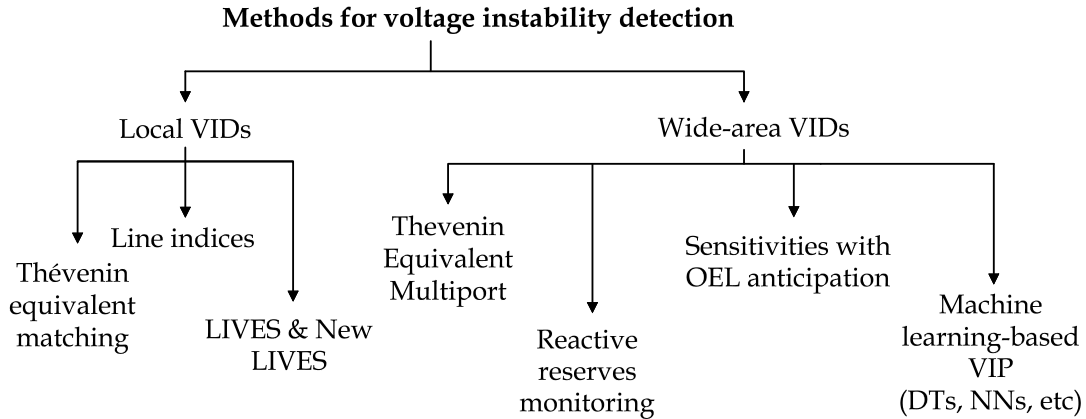


Figure 2.5: An overview of different types of methods aimed for VID.

2.3 An overview of VID methods

The aim of VID methods is to as fast as possible detect, or even predict, the onset of voltage instability. Numerous different methods have been proposed in the literature where the complexity and the requirements on the available measurement infrastructure vary significantly. The methods may also significantly differ in the accuracy and speed of the detection, as well as robustness to errors and other functionality. Available methods for VID can be divided into two different main categories:

1. *VID methods based on local measurements:* A VID method is in this thesis defined to be local if it relies on measurements from only two or fewer buses. Thus, those VID methods where measurements are required on both sides of, for instance, a transmission corridor, are also considered to be local VID methods.
2. *VID methods based on the observability of the whole region:* These methods are generally more accurate than the VSIs based on local measurements. However, as the name indicates, they require full or close to full observability of the monitored region and the measurements used in these models should preferably be filtered through a state estimator causing increased computation time and complexity.

The following section is a brief overview of the most prominent methods for VID. The number of different methods for VID is vast and the overview here is not meant to be exhaustive but is rather intended to provide an overview of the available methods and their advantages and disadvantages. More thorough overviews and comparisons for different types of VID methods can be found in [10, 39, 40] and in paper IV.

2.3.1 Thévenin equivalent matching

VID methods based on Thévenin equivalent (TE) matching use the TE impedance as an indicator of the margin to voltage instability [41]. Considering the simple two-bus system in Fig. 2.6, it can be shown that the maximum transferable apparent power in the system occurs when

$$|\bar{Z}_{th}| = |\bar{Z}_L| \quad (2.1)$$

where Z_{th} is TE impedance, and \bar{Z}_L is the complex load impedance. The relationship between the TE equivalent impedance and the other system parameters may be stated as:

$$\bar{E}_{th} = \bar{V}_L + \bar{Z}_{th} \cdot \bar{I} \quad (2.2)$$

where \bar{E}_{th} is the TE voltage, and \bar{V}_L and \bar{I} the load voltage and current, respectively. Using the relationship in (2.2), the values of \bar{Z}_{th} can be estimated. The real and imaginary values of \bar{E}_{th} and \bar{Z}_{th} in (2.2) results in four unknowns, requiring phasor measurements to be taken at two or more times to solve for the unknown parameters. By tracking and comparing the TE impedance and the load impedance, the system operator can assess the load margin to instability, for each bus in the system the index is computed for. An indication of instability occurs when the load impedance becomes lower than the TE impedance. It should be noted that this margin is *not* the same as the VSM or the DVSM presented in Section 5, as it only provides a margin to instability in the system with respect to a theoretic *load increase*, and not an $N-1$ security margin.

In [42], an improved method for TE impedance matching was proposed which could anticipate the activation of OELs to increase the speed of the VID method. However, a clear drawback of the TE matching methods is that they cannot incorporate and foresee the load restoration that follows after a disturbance in the system, which is often the main driver for instability. Studies have also shown that the TE matching methods have problems to predict voltage instability when applied to multi-load power systems [43].




Figure 2.6: A Two-bus Thévenin Equivalent Circuit

2.3.2 Line VID methods

Line VID methods are based on measurements being available on both sides of a two-port transmission line. The number of different line VID methods is vast, although many are based on the same concepts. Over time, several line VID methods based on the concept of maximal transferable power have been developed. These are similar to the TE-based methods for the bus VSIs, with the difference being that measurements are required in each end of a transmission line. An example of an early developed line VID was the transmission path stability index in [44]. Here, the maximum power transfer occurs when the voltage drop equals the load-side voltage, according to:

$$TPSI = \frac{V_s}{2} - (V_s - V_r \cos \delta) \quad (2.3)$$

where V_s and V_r indicates the sending and receiving end voltage, and δ is the voltage angle difference between the two nodes. Other line VSIs based on similar concepts are the voltage collapse proximity indicators in [45], where four indicators are developed, based on the maximum transferable power and the maximum possible line losses that may occur over a transmission line.

Other formulations of line VID methods include methods that are based on solutions to the classical voltage-power equation for a two-bus system with negligible resistance, previously presented in (5.1). It can be shown that the maximum power transfer occurs when the value of the inner square root in (5.1) is zero. A line index called L_{mn} based on this formulation is presented in [46]. The index can be formulated as:

$$\frac{4XQ_r}{[V_s \sin(\theta - \delta)]^2} = L_{mn} \leq 1.0 \quad (2.4)$$

where θ is the impedance angle and X is the line reactance. Instability is indicated whenever the stability index L_{mn} exceeds a value of 1.

2.3.3 LIVES and new LIVES

In [47, 48], a method called local identification of voltage emergency situations (LIVES) is introduced and tested. The LIVES stability criterion is based on monitoring the change in the secondary voltage after a tap decrease on the primary side ($\Delta r < 0$) of a LTC transformer, which simplified may be stated as:

$$\frac{\Delta V_2}{\Delta r} < 0 \quad (2.5)$$

where ΔV_2 is the change in the secondary voltage. Thus, if a tap decrease leads to a negative change in ΔV_2 , this indicates an unstable condition. Further, the criterion indirectly takes into account the effect of other taps acting in the system

as it can observe the net effect of various LTCs over a cycle of tap operations. In [49], this concept is developed further by monitoring the stability condition of (2.5), solely from the transformer bus, by assuming that primary voltage and current measurements are available. The decreasing tap change is measured indirectly as a conductance increase seen from the primary side, whilst the secondary voltage is indirectly monitored as an increase of consumed active power, P . The new index, denoted as the New LIVES Index (NLI) is formulated as:

$$NLI = \frac{\Delta P}{\Delta G_1} > 0 \quad (2.6)$$

where

$$G_1 = \text{Re}\{\bar{I}_1/\bar{V}_1\}$$

Simulations show promise during several different grid conditions and topologies, allowing early indication of impending voltage collapses. The method is further tested in [50], where the method is extended and applied for distance relays of transmission lines feeding weak areas.

2.3.4 Thévenin equivalent coupled single-port methods

In [43], it was shown that the TE matching methods do not work properly for multi-load systems. To overcome these difficulties, a concept called *coupled* single-port circuits was introduced, where an additional term modeling the coupling effects of generators and other load variations of other buses was added. In [51], it was found that the coupled single-port model may still yield underestimations if loads are not proportionally increasing. A modified coupled single-port model was then proposed to handle the underestimations that occur if loads are not proportionally increasing. Another extension was proposed in [52], which could better incorporate the dynamic nature of the grid equivalence in the estimations. Here, adjustments to the equivalent parameters of the coupled single-port model were calculated from two consecutive phasor measurements at the corresponding bus to capture the power system evolution.

It should be noted that these methods rely on PMUs covering *all* relevant generation and load buses, requiring a very well developed PMU configuration.

2.3.5 Reactive reserves monitoring

Voltage instability is highly related to the incapability of the system to provide sufficient reactive power reserves to maintain voltages around nominal values. By monitoring the reactive power reserves of system components such as synchronous generators and static var compensators, the remaining reserves can be used as an indicator for VID [10]. Several studies have examined reactive power reserves as an indicator for VID, examples including [53–56]. However, a difficulty in using reactive

power reserves for VID is that reactive power cannot be effectively transmitted over long distances on inductive transmission lines. Hence, although reactive reserves do exist in certain areas of a power system, it is infeasible to transmit it to the areas with low system voltages. Thus, methods using reactive reserves monitoring need to only take into account that reactive power reserves need to be close to the affected area to be effective [54, 55].

2.3.6 Sensitivities with OEL anticipation

In [18], a method using the sensitivities of reactive power generation to reactive power loads is considered. The method fits an extended set of algebraic equations to the sampled state of a power system, which is either gathered directly from wide-area phasor measurements or from a state estimator. Then, a sensitivity computation is computed to identify when a combination of load powers has passed through the system's maximum. The method does not require explicit modeling of the system loads, but will take into account OEL activation, either directly from measurements or by anticipation techniques. The method was tested and compared to conventional TE matching methods in [57], where the method proved to be significantly faster than conventional TE methods to identify voltage instability.

2.3.7 Machine learning-based VIP

The concept of voltage instability prediction (VIP) significantly differs from conventional VID methods. Most methods for VID attempts to detect when the system is close, or have already reached, the point of maximum load power. However, when this point is reached, instability can develop quickly and the remaining time for system operators to initiate emergency control actions may be too short. In contrast, machine learning-based methods for VIP attempts to *predict* the onset of instability. This would allow system operators to get an indication, almost instantaneously after a disturbance, whether that disturbance is going to cause a voltage collapse in the near future. These methods are generally based on off-line training of a chosen ML algorithm to be able to correlate a certain state space with the future state of the system. In Chapter 7, a new method for VIP is proposed and presented.

CHAPTER 3

Neural networks and machine learning

This chapter provides a theoretical background of machine learning and supervised learning. The general structure and training aspects of neural networks are presented. Furthermore, an introduction of recurrent neural networks and the long-short term memory architecture is presented. This theory serves as an overview of the subject which is the main method used in the following chapters.

3.1 Machine learning: an overview

Machine learning (ML) can be defined as a set of methods and statistical models used to perform specific tasks without using explicit instructions; rather relying on patterns and inference on data. ML algorithms differ in their approach, the type of task or problem that they are intended to solve, and the type of data used as input and output in training. ML is often divided into three main types of learning, namely: supervised learning, unsupervised learning, and reinforcement learning. Unsupervised learning is used to find hidden patterns and structures in data, while reinforcement learning refers to algorithms used to perform optimal-based decisions, often under various degrees of uncertainty. In supervised learning, which is the approach mainly used in this thesis, the ML algorithm is trained on a set of data that contains *both* the inputs and the desired outputs, where typical applications include classification or regression.

One of the simplest supervised learning algorithms is linear regression, where we are given a set of N input-output training data pairs $\{(\mathbf{x}_1, \mathbf{y}_1) \dots (\mathbf{x}_N, \mathbf{y}_N)\}$. A linear regression model have the general form of $f(\mathbf{x}) = \mathbf{x}\mathbf{W} + \mathbf{b}$, which is used to generate a linear function mapping each \mathbf{x}_i to each \mathbf{y}_i . In training of the linear regression model, the aim is to find the most suitable parameters for the weight matrix \mathbf{W} and the bias vector \mathbf{b} , to minimize a chosen objective function (commonly the average squared error). Once the algorithm is fully trained, it is capable to predict or

estimate similar outputs from now unseen input data. However, in more general cases where the relation between \mathbf{x} and \mathbf{y} is not necessarily linear, we may need to develop a nonlinear function mapping the inputs to the outputs [58]. In this thesis, we will mainly rely on various methods based on *neural networks* to develop these nonlinear functions.

3.2 Feed-forward neural networks

Feedforward neural networks (NNs), also known as multilayer perceptrons, are loosely inspired by the neurons in the human brain and its ability to classify and learn events from input data. The strength of NNs lies both in their capability of learning and approximating nonlinear functions and the scaling performance of the methods when trained on large sets of data. The universal approximation theorem, a famous theorem in NN mathematics, states that a feed-forward NN with a single hidden layer can approximate *any* given function arbitrarily well, provided that sufficiently many neurons are available in the hidden layer [59].

A typical NN with a single layer of hidden units is presented in Fig. 3.1. The NN consists of connected nodes, known as artificial neurons, stacked in different layers. Neurons in one layer only connect to neurons of the immediately preceding and the immediately following layers. The input layer receives external data, while the layer that produces the final result is denoted as the output layer. Between these layers are the hidden layers, where each neuron has a nonlinear activation function, which imitates the actions of synapses in a biological brain. The connections between each neuron is represented by weight parameters and the aim in training a NN is to adjust these parameters so that the network can accurately map the inputs to the outputs. In the following sections, we provide a general overview of the steps needed in training the network. We provide the NN model for a single hidden layer, which can then be generalized to NN models using multiple hidden layers.




Figure 3.1: A neural network with a single hidden layer.

3.2.1 The forward pass

The *forward pass* in a NN is used to compute the outputs, both in training as a means to adjust the weight parameters, but also in actual implementations where the outputs of the network are used as estimates or predictions. Following [58], the forward pass of the single hidden layered NN is presented below. Superscripts M and N refer to the number of input features and the number of hidden neuron cells in the hidden layer. Vector notation is used, meaning that, for instance, the output of the hidden layer $\mathbf{h}(\mathbf{x})$ is not the output of a single neuron cell, but the output of a vector of N neuron cells. The operation is summarized here for a single sample by the following steps.

The input layer first passes a row vector $\mathbf{x} \in \mathbb{R}^M$ of inputs through the weight matrix $\mathbf{W}_1 \in \mathbb{R}^{N \times M}$, illustrated by the lines connecting each of the cells in Fig. 3.1. The outputs of the hidden layer is computed by applying an element-wise nonlinear function on the weighted input values:

$$\mathbf{h}(\mathbf{x}) = \sigma(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) \quad (3.1)$$

where σ is a nonlinear activation function, and $\mathbf{b}_1 \in \mathbb{R}^N$ is a bias term for the first weight matrix. Common activation functions used in the hidden layers of NNs are the hyperbolic tangent function (Tanh) or the rectified linear (ReLU) function. The outputs of the hidden layer are passed through the second weight matrix \mathbf{W}_2 , which is used to compute the outputs \mathbf{z} :

$$\mathbf{z} = \mathbf{W}_2 \mathbf{h} + \mathbf{b}_2 \quad (3.2)$$

where \mathbf{b}_2 is a bias term for the second weight matrix. The dimensions of \mathbf{W}_2 , \mathbf{b}_2 , and consequently \mathbf{z} depends on the application of the NN and the number of target values. A final activation function is applied in the output layer to generate the estimated target values: $\hat{\mathbf{y}} = f(\mathbf{z})$. For binary classification, a sigmoid activation function is applied on the sum estimated outputs, generating \hat{y} ranging from 0 to 1:

$$\hat{y} = \frac{1}{1 + e^{-z}} \quad (3.3)$$

In multiclass classification, a softmax activation function is applied that normalizes it into a probability distribution consisting of K number of probabilities proportional to the exponents of the input numbers:

$$\hat{\mathbf{y}} = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}} \quad (3.4)$$

for $i = 1, \dots, K$ and $\mathbf{z} = (z_1, \dots, z_K) \in \mathbb{R}^K$, where K is the number of observed target values. In regression, \hat{y} is simply a linear combination of z .

3.2.2 Loss and learning

The aim of training a NN is to tune the weight matrices connecting each layer of neurons so that the network can accurately model the relationship between the

inputs and the target values. A suitable loss function $L(\hat{\mathbf{y}}, \mathbf{y})$ is first applied to the difference between the estimated target vectors and the true target vectors that are used in the training set. To use a NN for classification, cross-entropy (log-loss) functions are commonly used, while for regression purposes other means of loss such as the mean squared error (MSE) are applied.

The tuning of the parameters in the weight matrices is performed iteratively using gradient-based optimization algorithms, which refers to moving along an error gradient towards some minimum level of error of a defined objective function J . Using the whole training set to update the weight parameters is commonly referred to as batch gradient descent, in which the objective function can be defined as:

$$J = \frac{1}{S} \sum_{i=1}^S L(\hat{\mathbf{y}}, \mathbf{y}) \quad (3.5)$$

where S is the total number of samples in the training set. In batch gradient descent, the objective function is defined as the average loss over the *whole* training set. However, batch gradient descent is often computationally inefficient and requires significant memory when a network is trained on large sets of data. To overcome this issue, training methods that use smaller subsets of the total training set is often used. In *mini-batch gradient descent*, a smaller batch is used to update the parameters. When all batches in the whole training set have been used to update the weight parameters, it is said that the network has been trained for one *epoch*.

The gradient of the objective function is computed with respect to the weight parameters and the bias vectors. A method known as *backpropagation*, where the computed gradient is passed back along each layer to update the weight parameters connecting each layer. The weight parameters are updated iteratively in small steps, and it is common to train a NN for several epochs to reach good performance on the training set.

3.2.3 Overfitting and validation

The aim in training a NN is not to fit the data on the *specific* training set, but rather to generalize the training so that the network can provide good mapping on yet unseen data. A common problem when training NNs is overfitting, where the network is trained too closely to the specific training set, with the result that it fails to fit additional data or predict future observations accurately.

To ensure that the network is not overfitting on the training data, the performance of the network on a separate validation set is commonly monitored. In case the validation loss starts to increase, the training of the network can either be stopped, or various regularization techniques can be applied to avoid overfitting. Many regularization techniques are based on limiting the capacity of the networks, by adding different penalties to the objective function J . A very common and simple regularization function is the $L2$ parameter norm penalty, also known as weight decay. This technique penalizes large numbers of the weight parameters by adding a regularization term $\Omega = \lambda \frac{1}{2m} \sum_{i=1}^k \|\mathbf{w}\|_i^2$ to the objective function. The value of λ controls

how strict the regularization is, and m is the number of training samples. A too small value of λ may result in an insufficient regularization with overfitting as a result, while a too large value may result in *underfitting* or high variance, with poor performance on both the training set and the validation set. Another popular regularization technique is called *dropout*, where a certain percentage of the connections between each layer are randomly masked (or "dropped"). This technique ensures that the network does not rely too heavily on certain connections.

3.2.4 Hyperparameters and network depth

Hyperparameters are parameters that control the training of the NN, and include parameters such as the learning rate of the optimization algorithm, the number of epochs the network is trained, or the structure and number of neurons in the hidden layers. The hyperparameters play a crucial part in the performance of the NN and a hyperparameter search is generally conducted where different combinations of hyperparameters are assessed and tested [58].

Deeper network architectures, with several hidden layers ordered in chain-like structures, are often able to use far fewer hidden units per layer, and thus far fewer parameters, compared to more shallow networks used in the same applications. Deeper architectures are also found to better generalize the performance on unseen data. However, the drawback of deeper architectures is that these also tend to be harder to optimize [58].

3.3 Recurrent neural networks

In many applications, the capability to provide accurate classifications or estimations at a specific time t would be more accurate if it was possible to account for previous or historic data. Recurrent neural networks (RNNs) provides this capability by using sequence-based networks that are adapted for processing sequences of input data, capable of utilizing both current and past data [60]. The concept of RNNs is not new, but has received an increasing amount of attention in recent years. In power systems, RNNs have previously been used in applications ranging from market forecasting [15], transient stability assessment [16], and in power quality assessments [17].

A typical RNN-sequence is illustrated in Fig. 3.2, where each block has a directed connection to the following block in the sequence. The main difference from a feedforward NN is that the output of each block depends not only on the input vector at the current time step but also on the output from previous blocks in the sequence. If the block is the first one in the sequence, the inputs are made up solely by the input vector. Depending on the architecture of the RNN, the output vector of each block can be used for both classification at the current time step, and/or used as an input to the following block. Each block contains interior connections, weights,




Figure 3.2: A sequence of consecutive RNN blocks.

and non-linear activation functions. The block's complexity ranges from those of simple RNNs using conventional neuron layers, to more advanced structures such as the long-short term memory (LSTM). RNN are generally trained using an approach called *backpropagation-through-time*. Backpropagation-through-time is similar to the backpropagation used in conventional NN, with the difference that the error gradients also have to be propagated back in time through the RNN-sequence [61].

3.3.1 Long short-term memory

The standard implementation of RNN has difficulties in capturing long-term dependencies of events that are significantly separated in time. When the error signal is passed back through many RNN sequences, it tends to either blow up or vanish [62]. This effect is generally recognized as *vanishing* or *exploding gradients*. In an LSTM network, the information and the error gradients can be propagated through time within an internal state memory cell, making the network capable of memorizing features of significance over time [63].

A typical LSTM-block is illustrated in Fig 3.3. The state memory cell, illustrated by the light grey area, is controlled by nonlinear gating units that regulate the flow in and out of the cell [64]. Following [63] and [64], the forward operation of an LSTM block is summarized below. It should be noted that each block consists of a number of hidden LSTM cells. Vector notation is used, meaning that, for instance, the hidden state vector \mathbf{h}^t is not the output of a single LSTM-cell at time t , but the output of a vector of N LSTM-cells. The operation of an LSTM block at a time t may then be summarized by:

$$\mathbf{f}^t = \sigma(\mathbf{W}_f \mathbf{x}^t + \mathbf{U}_f \mathbf{h}^{t-1} + \mathbf{b}_f) \quad (3.6)$$

$$\mathbf{i}^t = \sigma(\mathbf{W}_i \mathbf{x}^t + \mathbf{U}_i \mathbf{h}^{t-1} + \mathbf{b}_i) \quad (3.7)$$

$$\tilde{\mathbf{c}}^t = \tanh(\mathbf{W}_c \mathbf{x}^t + \mathbf{U}_c \mathbf{h}^{t-1} + \mathbf{b}_c) \quad (3.8)$$

$$\mathbf{c}^t = \mathbf{f}^t \odot \mathbf{c}^{t-1} + \mathbf{i}^t \odot \tilde{\mathbf{c}}^t \quad (3.9)$$

$$\mathbf{o}^t = \sigma(\mathbf{W}_o \mathbf{x}^t + \mathbf{U}_o \mathbf{h}^{t-1} + \mathbf{b}_o) \quad (3.10)$$

$$\mathbf{h}^t = \mathbf{o}^t \odot \tanh(\mathbf{c}^t), \quad (3.11)$$

where element-wise multiplication is denoted by \odot , σ is the logistic sigmoid function, \tanh is the hyperbolic tangent function, and with the following variables:

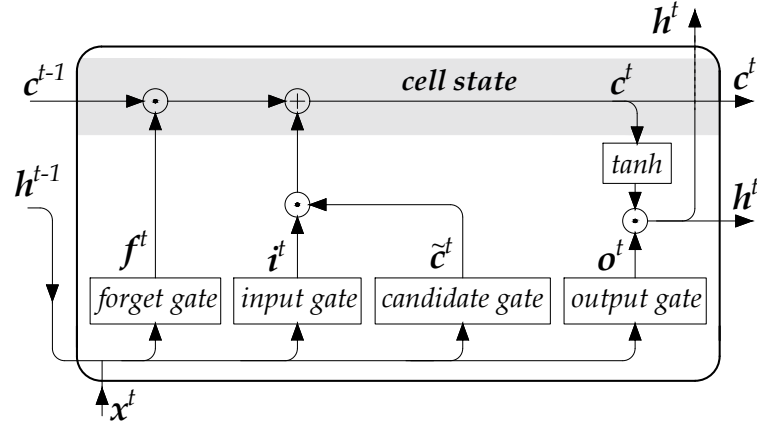


Figure 3.3: Detailed schematics of an LSTM block.

- $\mathbf{x}^t \in \mathbb{R}^M$: input vector to an LSTM block
- $\mathbf{h}^t, \mathbf{h}^{t-1} \in \mathbb{R}^N$: output vector at time t respectively $t-1$
- $\mathbf{f}^t \in \mathbb{R}^N$: activation vector of the forget gate
- $\mathbf{i}^t \in \mathbb{R}^N$: activation vector of the input gate
- $\tilde{\mathbf{c}}^t \in \mathbb{R}^N$: vector of the candidate gate
- $\mathbf{c}^t \in \mathbb{R}^N$: cell state memory vector
- $\mathbf{o}^t \in \mathbb{R}^N$: activation vector of the output gate

where \mathbf{W} , \mathbf{U} , and \mathbf{b} represent the weight matrices and bias vectors for each gate. The superscripts M and N refer to the number of inputs and hidden LSTM cells in each LSTM block, respectively.

The information stored in the state memory cell is regulated by the operation of the different gates, as illustrated in Fig. 3.3. By the operation of (3.6), the forget gate controls what information should be stored from the previous memory cell state, and what can be discarded as irrelevant. The input gate and candidate gate control and update the memory cell state with new information by the operation of (3.7)–(3.8). In (3.9), the state memory cell is first updated by an element-wise multiplication of the previous cell state memory vector and the resulting vector of the forget gate. Then, the state memory cell is updated with new values provided by an element-wise multiplication of the resulting vectors from the input gate and the candidate gate. Equations (3.10)–(3.11) show how the hidden state is updated by the operation of the output gate, modulated by the updated cell state memory vector.

An LSTM *network* may then be constructed by creating a sequence of several LSTM blocks. A partition of an LSTM sequence is illustrated in Fig. 3.4, where each block has a directed connection to the following block in the sequence. If the block is the first one in the sequence, the past system state is initialized with a preset value.

For a deep LSTM network, with several stacked layers, the inputs to the deeper layers consist of the hidden states of LSTM blocks of previous layers. The cell state memory is only passed along the time sequence between LSTM blocks of the same layer. Typically, for classification purposes, an output vector \mathbf{y} is generated by applying a nonlinear function of the hidden state implemented by a separate feedforward NN. Depending on the application of the network, output vectors may be computed for a single, or several, LSTM block's hidden states.




Figure 3.4: An LSTM sequence with a directed connection between the blocks.

CHAPTER 4

Description of a real-time voltage stability assessment tool

In this chapter, the overall functionality of the proposed real-time voltage stability assessment tool (RVSAT) is presented. The RVSAT is based on two developed methods presented in subsequent chapters. The first method is developed for preventive purposes, where a method for fast DVSM estimation is developed. The second method is developed for emergency purposes, where a method for VIP based on an LSTM network is developed. In Chapters 6 and 7, each of the two developed methods will be presented in more depth. Furthermore, the test system used in evaluating the developed RVSAT is presented and its general characteristics are briefly discussed.

4.1 Real-time voltage stability assessment tool

This thesis aims to develop a tool for real-time voltage stability assessment (RVSAT) that can support system operators and increase the transmission capacity of a power system. An overview of the proposed RVSAT in the setting of a monitoring system in a power system is presented in Fig. 4.1. The functionality of the RVSAT is divided to handle the two complementary tasks in VSA: i) *emergency* and ii) *preventive* monitoring.

The proposed method for emergency monitoring is based on using a RNN with LSTM for VIP, from now on abbreviated into LSTM-VIP. The method is designed to take current and historic measurements to assess whether the *current* state will cause voltage stability issues several minutes into the future. As time progresses and if new events occur in the system, the network updates the assessment continuously. Stability warnings are then passed to the system operator, which can initiate emergency control actions, such as load shedding in certain areas. The network is also adapted to be able to indicate where in the system instability emerges, following the approach developed in [15], allowing more cost-effective countermeasures.

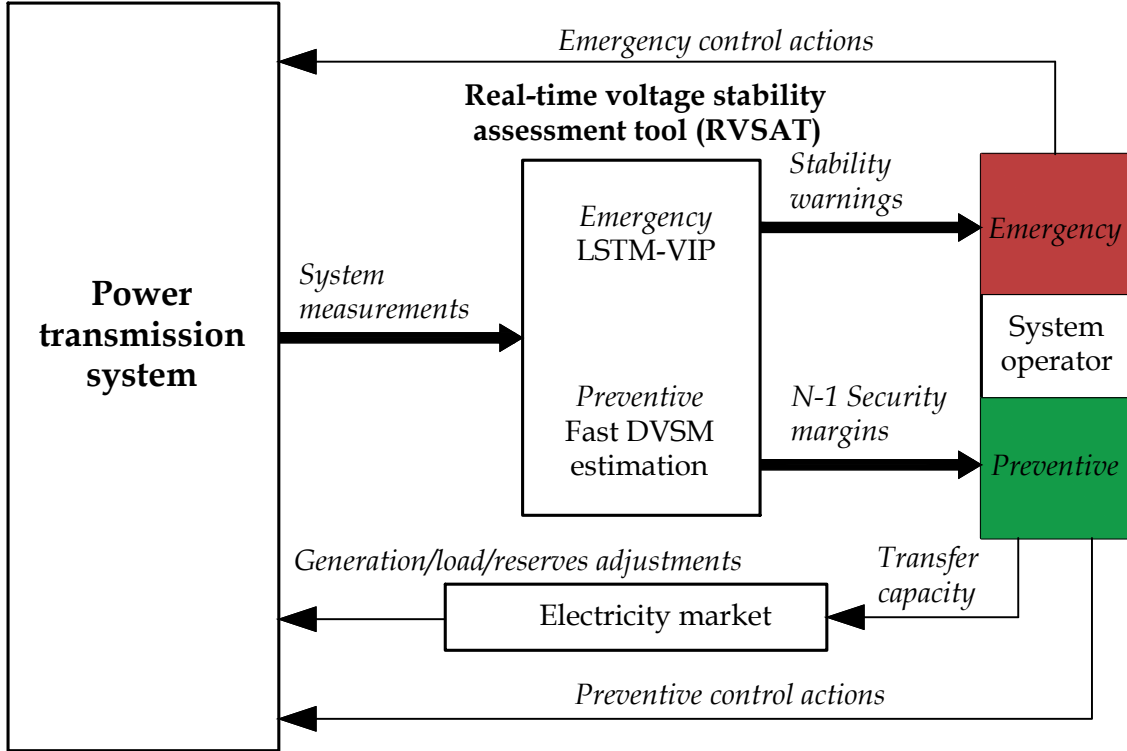


Figure 4.1: An overview of the proposed RVSAT and the signals and control actions in a monitoring system.

The proposed method for preventive monitoring is based on using conventional NNs to increase the computational efficiency in estimating the DVSM. The method computes the $N-1$ voltage security margins which are then passed to the system operator. The security margins are used to compute the transfer capacity in the system which is used in up and down regulation. Through price signals, such as price area differences, the actors on the market can adapt to the changing system transmission capacity. The system operator can also initiate preventive control actions, for instance by controlling reactive power components, to adapt to the estimated security margins.

Both methods use real-time system measurements gathered directly from the power transmission system. The requirements on the system measurements, such as measurement update rates and the availability of phasor measurements, are discussed in each of the chapters covering the functionality of each method. To assure that errors and missing values are filtered out, measurements are assumed to always be preceded by a state estimator.

4.2 Test system

The methods are tested on an updated version of the classical Nordic32 test system, presented in [65]. The system is fictitious but is similar to the Swedish and the Nordic power system. A one-line diagram of the test system is presented in Fig. 4.2. The system is divided into four different regions:

- **"North"**: mainly consists of hydro generation and some smaller loads.
- **"Central"**: the largest load center with significant generation of thermal power generation.
- **"Eq"**: An equivalent of an external system connected to the "North".
- **"South"**: An area with thermal generation which is loosely connected to the "Central" area.

The Nordic32 test system has long transmission lines of 400-kV and 220 kV nominal voltage. The test system also includes a representation of regional systems operating at 130 kV. Table 4.1 provides the active power generation and load in each area and for the whole test system.

The system is heavily loaded with large power transfers mainly between the areas "North" and "South". The transferable power is limited by the reactive power capabilities of generators in both of these areas. A disturbance in the transmission capacity connecting the "North" and the "Central" areas is critical for the stability of the system. A reduced transmission capacity, caused by, for instance, a disconnected transmission line, would increase the strain on the remaining lines in the system. The increased current in the remaining transmission lines would increase the reactive power losses in the system, which would cause lower system voltages. Voltage-dependent loads are restored through the actions of LTCs. The load restoration has an adverse effect on the voltage stability of the system as the restored loads will cause increased stress on the remaining lines in the system and may cause the system to deteriorate further. Reactive power limits of generators are enforced by

Table 4.1: Active power generation and load for the Nordic 32 test system.

Area	Generated power (MW)	Load (MW)
North	4628.5	1180.0
Central	2850.0	6190.0
South	1590.0	1390.0
Eq.	2437.4	2300.0
Total	11505.9	11060.0

OELs which further limits the capability to maintain nominal voltages in the system.

For the developed Nordic32 system, two developed operating points were presented in [65]. "Operating point A" is an unsecure operating point and an outage of either a larger thermal generating unit in the "Central" area or a disconnection of a transmission line connecting the "North" and "Central" areas may cause instability. "Operating point B" is a secure operating point which should be able to withstand the loss of any single component. In the following simulations, the test system is mainly operated in and around the secure "operating point B". All time-domain simulations have been performed using PSS®E 34.2.0 with its built-in dynamic models.

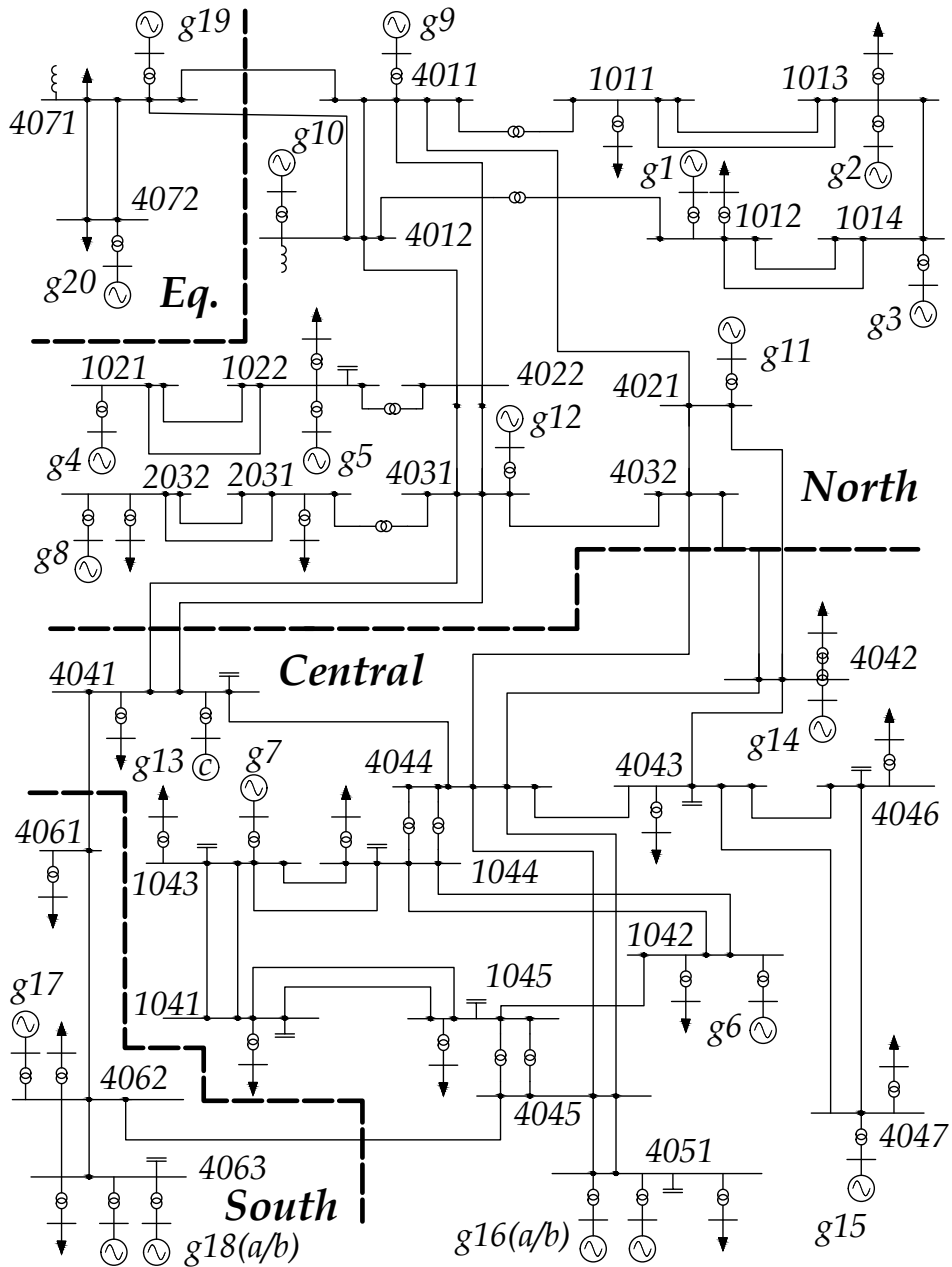


Figure 4.2: The Nordic32 test system used in testing the developed algorithms.

CHAPTER 5

Comparison of dynamic and static voltage security margins

In this chapter, the difference between static and dynamic voltage security is established. First, the theoretical difference between the two methods are established and then visualized using the concept of transient P - V curves. Second, numerical comparisons are provided to further examine the differences between the different security margins under various load configurations and types of disturbances. The chapter is mainly based on the results established in Paper I.

5.1 Theoretical comparison of DVSM and VSM

In this section, the theoretical difference between the DVSM and the conventional VSM is established and then visualized using the concept of transient P - V curves. A small test system's dynamic response following a disturbance is used in the analysis.

5.1.1 Small test system

In the following examples, the impact of load dynamics and the voltage control devices (e.g. excitation control for synchronous generators and synchronous condensers) are mainly taken into account in the analysis. The small 2-bus test system in Fig. 5.1 is used in the analysis. It consists of a controlled sending end voltage ($E\angle 0$), supplied by a voltage source through a reactance X_f . A complex load ($P + jQ$) is fed through a number of lines represented by inductances with the total reactance of X_t .

A popular method in static voltage stability analysis is to use P - V curves, where the receiving end voltage is plotted with respect to an increasing active load transfer in the system. In the following figures in this section, P - V curves for the case when $E = 1.05$ pu, $X_t = 0.4$ pu, and a fully active power load are illustrated.




Figure 5.1: Simple 2-bus system used in the analysis.

The reactance X_f is initially neglected, but will be introduced in Section 5.1.4. An additional P - V curve is plotted in each figure for a $N-1$ case when one line has been disconnected (increasing X_t to 0.5 pu). Assuming lossless transmission, the curves are developed using the classic voltage equation for a two bus system, given by [9]:

$$V = \sqrt{\frac{E^2}{2} - QX_t \pm \sqrt{\frac{E^4}{4} - X_t^2 P^2 - X_t E^2 Q}} \quad (5.1)$$

where the upper part of each P - V curve corresponds to the solution of (5.1) with the plus sign, while the lower part of each curve corresponds to the solution with the minus sign.

The voltage instability mechanism is mainly driven by loads and the impact of load modelling in voltage stability analysis is imperative [9]. The power consumption of loads are affected by the system voltages and different load models are often used to characterize this relationship. One conventional model is the *exponential load model*, which is given by:

$$P = zP_0 \left(\frac{V}{V_0} \right)^\alpha \quad (5.2)$$

$$Q = zQ_0 \left(\frac{V}{V_0} \right)^\beta \quad (5.3)$$

where P_0 and Q_0 are the active and reactive power consumed at voltage V equal to the reference voltage V_0 when $z = 1$. z is a dimensionless and independent variable indicating the actual loading of the system [9]. The voltage dependency is modeled by the α and β parameters, where $\alpha = \beta = \{0, 1, 2\}$ represents constant power (MVA), constant current, and constant impedance characteristics, respectively.

5.1.2 Estimating DVSM and VSM

The difference in estimation methods if the VSM and the DVSM is illustrated in Fig. 5.2 using the developed P - V curves. The security margin is defined as the change in loading from an initial operating condition (OC) to the $N-1$ collapse point. It should be noted that in real applications, the limit is often significantly




Figure 5.2: Difference between VSM and DVSM illustrated.

smaller due to the other stopping criteria such as too low system voltages. However, for better illustration purposes, the former limit is used.

In static VSM estimation, the initial post-contingency operating point is found by first introducing a contingency on the initial OC, which is followed by solving the resulting power flow study. This is illustrated in Fig. 5.2 by moving along arrow **1'**. The stability limit is then traced along the solution path by iteratively increasing the system stress until the critical point is reached, moving along the arrow **2'**. Continuation power flow methods are preferably used to avoid convergence problems close to the collapse point [12, 13]. The distance between the pre-contingency operating point and the $N-1$ critical point constitutes the VSM.

The steps of estimating the DVSM are conceptually different from VSM, where instead the dynamic security of the system is being tested with an increasing stress level in the system, illustrated by arrow **1** in Fig. 5.2 [11]. For every new pre-contingency operating point (an increase in system stress), a time-domain simulation is initiated where the system response following a disturbance is studied. The simulation runs until the system stabilizes or becomes unstable. The final pre-contingency operating point that is tested and still provides a stable operating point is illustrated by moving along the curved arrow **2** in Fig. 5.2. The distance between the initial operating point and the last pre-contingency operating point that can still handle a dimensioning contingency without causing a voltage collapse, represents the DVSM.

5.1.3 Load response after a disturbance

In static VSM, loads are often recognized to maintain a constant MVA characteristic. This assumption is often true from a long-term *system* perspective, but does not necessarily mean that the loads themselves behave as constant MVA loads. Equipment




Figure 5.3: Example of slow load restoration after a contingency.

and control mechanisms such as load tap changers (LTCs) and voltage regulators will restore load voltages following a disturbance, resulting in recovered load levels even for loads with constant impedance characteristics [66]. However, even though loads are considered to have long-term MVA characteristics, they do not necessarily behave as static MVA loads following a disturbance. Assuming a sudden voltage change, the loads will initially change according to their instantaneous characteristics, for instance, a mix of constant impedance and constant current load [66]. Then, they will adjust their impedance or the drawn current to restore the load to their original level.

This load restoration event following a disturbance, tripping of a line in the system, is illustrated in Fig. 5.3. The initial OC is located at **A**. Instantly after a disturbance, the load is assumed to have constant impedance characteristics, which results in a change in operating point from **A** to **B**. Load restoration dynamics then change the operating point from **B** to **C**, which corresponds to the same initial load level as point **A**.

5.1.4 Transient P-V curves and fast load dynamics

In [67], it was shown that if the system starts at a stable equilibrium and is slowly stressed towards the collapse point without encountering oscillations or other limit-induced events (e.g. reactive power limits for generators), the static equations are sufficient to locate the exact collapse point experienced by the dynamic system. However, the majority of voltage collapse incidents experienced so far have resulted from large disturbances, typically by the loss of generation or transmission capacity [11]. In static VSM estimation, the transient state of such events can be neglected using the assumptions (i) that loads do not behave as constant MVA loads just after a disturbance, and (ii) that load dynamics acts significantly slower compared to the voltage control dynamics of, for instance, excitation system of generators and synchronous condensers [66]. Hence, the transient impact of voltage control

dynamics can be neglected and the assumptions developed in [67] would still be valid.

However, load dynamics of induction motors and power electronic loads, such as chargers for electric vehicles, are inherently fast. For these components, the load is often restored in a time frame within a second, similar to that of most excitation systems [9, 68], causing the assumptions used in conventional VSM estimation to falter. In [66], a concept called transient P - V curves was adopted to visualize the dynamic impact of voltage control on the static P - V curves. Here, the same approach is used when the difference between VSM and DVSM is illustrated. The transient P - V curves can be obtained by modelling and taking into account the dynamic impact of having the voltage source in Fig. 5.1 behind the reactance X_f . The assumption used in conventional VSM estimation, that excitation control instantly will restore E to its pre-contingency value after a disturbance, will thus no longer be true. Instead, E will initially be affected by events in the system but is controlled back to its nominal value by excitation control of the voltage source. As the main purpose here is to provide a principal understanding of the concept, the transient P - V curves in the following figures are hypothetical. Similarly, the curves illustrating the *fast* load restoration dynamics from a constant impedance load to a constant power load are drawn to allow better understanding of the concept.

In Fig. 5.4, the dynamic response following a disturbance is illustrated for a secure initial OC. The transient P - V curves and the load restoration curves are illustrated using different shades of grey, where a lighter shade indicates closer in time to the disturbance. The time just after a disturbance is indicated by t_1 , while t_3 relates to the time when all short-term dynamics have already taken place. The load is assumed to have long-term constant MVA characteristic, but just after a disturbance, the load will initially change to a constant impedance characteristic. Then, by fast load restoration, the load is quickly restored to a constant MVA characteristic.

The initial OC is found in point **A**. Just after a disturbance (at t_1), the bus voltages drop caused by a larger current being transmitted through the remaining lines. The reactive power losses increases in the system, and a larger current is being transmitted through the reactance X_f found in in Fig. 5.1. The larger current causes the voltage E to drop initially, resulting in the P - V curve being shifted to the left (the lightly shaded P - V curve). As a result of the initial load characteristics and the shifted P - V curve, the operating point moves along the arrow to operating point **B**. After the shifted operating point, two separate dynamic responses are initiated. The voltage control dynamics, here illustrating the excitation system response for a synchronous generator, is restoring the terminal voltage E to its nominal value. This causes the transient P - V curve to shift back towards the P - V curve for the static $N-1$ case. Simultaneously, the fast load dynamics is restoring the load from the initial load characteristics back to a constant MVA characteristic. As an effect of the voltage control dynamics and the load restoration, the operating point moves along the arrows from **B** to **C**, then finally from **C** to **D**. In this case, the system was found to be stable even after the disturbance with the new operating point **D**.




Figure 5.4: Transient P - V curves for a secure initial operating condition.

In Fig. 5.5, the same system is slightly more stressed, with a higher level of initial transferred power. Just after the contingency, the operating point moves along the arrow from A' to B' , by same the reasoning as in the previous case. However, due to the fast load dynamics, there exists no intersection between the curves at t_2 , and without any emergency control actions, the system stability would be lost. The example in Fig. 5.5 illustrates a type of event that the DVSM could identify and take into account, which is not possible using a static VSM. It should be noted that the P - V curve for the $N-1$ case and the load characteristic at t_3 still intersect in this case, indicating that a static VSM would still classify the initial OC as secure.

5.1.5 DVSM versus VSM

The analysis in the previous section showed that the DVSM is to prefer over the static VSM in power systems with a large share of loads with fast restoration dynamics. Furthermore, the closer power systems are being operated to the limits of operation, the event illustrated in Fig. 5.5, the more likely it is that the system will become unstable during the transient state after a disturbance. However, the advantages of using DVSM is not limited solely to the short time instance after a larger disturbance has occurred in the system. The same type of events may occur significantly later in a voltage instability event, triggered by larger drops in system voltages from, for instance, OELs, or undervoltage tripping of generators. These types of events are generally referred to as short-term instability events *induced* by long-term dynamics [9]. It should be noted that methods based on QSS and combinations of QSS and time-domain simulations as was suggested in [24], cannot deal with those type of events.

A clear advantage of using the DVSM (and the method based on QSS), is that in static VSM, the notion of time is fully ignored, and by that the impact of, for in-




Figure 5.5: Transient P - V curves for an unsecure initial operating condition.

stance, timer settings of OELs, LTCs, and switched reactive power components [16]. Furthermore, equipment such as air conditioners, induction motors, and undervoltage relays, may either stall or trip due to temporary low voltages, which is an effect that can better be taken into account in either time-domain simulations or QSS simulations.

5.2 Numerical comparison of DVSM and VSM

In this section, a numerical comparison between the VSM and the DVSM is presented to further examine what parameters will cause the two methods to compute the security margins to differ. The Nordic 32 test system is used in the analysis.

5.2.1 Methodology for numerical comparison

The security margins are computed for two different types of disturbances, and under a large range of different load configurations. Only different configurations of the exponential load model have been taken into consideration. The two types of disturbances used in the simulations are the following:

- **Case A:** A three-phased fault for 100 milliseconds, followed by tripping the faulted line. The faulted line is the one connecting the two areas "North" and "Central" between bus 4032 to bus 4044.
- **Case B:** The exact same fault type as for Case A, with the difference that the fault is now cleared after 40 milliseconds by tripping the faulted line.

The fault clearing-time of 100 milliseconds is the same as the one used in [65]. The actual fault clearing-time is dependent on relay time (that is, fault detection and

relay logic's) and circuit breakers' operating time.

The security margins were computed by increasing the power transfer between the areas "North" and "Central". All loads in the "Central" area were increased incrementally and uniformly, starting from the secure "operating point B" in [65]. The increase in loading was met by simultaneously and uniformly increasing the generation of the generators in the "North" area. The distribution of the added load and generation was based on the initial load or the rated capacity of each generator. Thus, a bus with a larger initial load, or a generator with a higher rated capacity, received a larger share of the increased load and generation. The increased losses caused by a higher power transfer on the lines were assumed to be compensated either by a change in the output of the slack bus or by automatic changes in the governor models connected to each generator. All generation that could not be supplied by the regular generators was distributed to the slack bus generator in the system, g20.

Two different methods to compute the VSM are compared to the DVSM. As the tested methods to compute the security margins differ in their approach, some slight adjustments to the test system were required to ensure that the methods could be compared. The governor models of the hydro stations in the Nordic32 system were thus changed to the more generic governor model of IEESGO [69]. The new governor model, with its dynamical properties, is presented in Appendix A. The IEESGO model was chosen as it uses mechanical power as a direct reference for the controller which allows direct adjustments of generator power setpoints. The different methods to compute the security margins are the following:

1. **VSM based on full Newton-Raphson solutions (VSM_{FN}):** The static voltage security margin is preferably computed using methods stable close to the system critical point, such as the continuation power flow algorithm (CPF). However, it is not uncommon for the security margin to still be computed using conventional load flow computations.

The fault is first introduced to the base case. The power transfer is then increased in small increments of 10 MW, where for each increment of system loading, a new Newton-Raphson load flow is solved to estimate if the system is still stable. The system is considered unstable if either i) the load flow solution does not converge, or ii) any transmission bus voltage is lower than 0.9 pu. The Newton-Raphson load flow solutions were calculated assuming stepping tap adjustments, using the previous cases voltage magnitudes and angles as starting values (no flat start) and a maximum of 20 iterations.

2. **VSM based on dynamic loadability limits (VSM_{DLL}):** In [65], an alternative method to the CPF was used to compute the VSM, based on stressing a system in its post-contingency state using dynamic simulations. Here, the VSM is evaluated from time-domain simulations by *slowly ramping* up the system stress, allowing the system to stabilize between each system stress increment. The steps for the computation are the following:

- **Apply disturbance:** The contingency is applied in the system from the base case. The system is then allowed to stabilize by running the dynamic simulation without any additional disturbances for 300 seconds.
 - **Increase system stress:** The system load is then increased in increments of 5 MW, which is distributed among all the loads in the "Central" area. Simultaneously, the governor reference for all the generation units in the "North" area is increased to compensate for the increased load level. The new IEESGO model for the governors takes a mechanical power reference as inputs.
 - **Allow system to stabilize and check for stability:** After the increased system stress, the time-domain simulation runs for 100 seconds. The system is considered unstable if either the simulation crashes or if any transmission bus voltage is lower than 0.9 pu.
 - **Re-iterate:** The system stress increase is re-iterated until the system becomes unstable. The difference in loading from the base case to the final stable operating point in the stressed post-contingency system represents the computed VSM.
3. **DVSM based on time-domain simulations (DVSM):** The DVSM is estimated in the same approach as is later described in Section 6.1. To summarize the approach, the system stress is iteratively increased in the system's pre-contingency state. For each increment of system stress, a time-domain simulation is initiated where the system response following a disturbance is studied. The simulation runs until the system stabilizes or becomes unstable. The difference in loading between the initial operating point and the last pre-contingency operating point that can still handle a dimensioning contingency without causing a voltage collapse, represents the DVSM. A precision value of $\epsilon = 10MW$ was used in the simulations.

5.2.2 Results and discussion

In Table 5.1 and Table 5.2, the computed security margins for the different load configurations are presented for Case A and Case B. The largest difference between the three methods is found for cases with a high share of constant MVA characteristic of the active part of the loads. For these cases, the DVSM is significantly lower than the two measures of the VSM. The largest difference is found for Case 1A when the active part of the load has a fully constant power characteristic. Here, the DVSM is only 70 MW, while the VSM_{DLL} is 350 MW, and the VSM_{FN} is 298 MW. This is in accordance with the theory developed in Section 5.1.2, where a larger penetration of loads with fast load restoration (implicitly modeled by a constant MVA characteristic) will cause the DVSM to be lower than the VSM. Similarly, the DVSM and the VSM_{DLL} is found to be very close for system load configurations where the share of loads with constant MVA characteristic is fairly small, such as the cases 11-15 (both A and B).

Table 5.1: Case A: Three-phased fault during 0.1 seconds cleared by tripping line.

Case number	<i>Constant</i>			VSM_{FN} [MW]	VSM_{DLL} [MW]	DVSM [MW]
	MVA (P/Q) [%]	I (P/Q) [%]	Z (P/Q) [%]			
1A	100/0	0/0	0/100	298	350	70
2A	99/0	1/0	0/100	298	360	90
3A	95/0	5/0	0/100	298	380	260
4A	90/0	10/0	0/100	299	390	360
5A	80/0	20/0	0/100	298	400	410
6A	50/0	50/0	0/100	299	435	420
7A	100	0/50	0/50	299	330	80
8A	95/0	5/50	0/50	299	350	230
9A	90/0	10/50	0/50	299	380	330
10A	50/50	50/50	0/50	299	425	330
11A	0/0	100/0	0/100	299	455	460
12A	0/0	50/0	50/100	299	480	470
13A	33.33/0	33.33/0	33.33/100	299	450	450
14A	0	20/0	80/100	299	490	510
15A	20/0	0/0	80/100	299	475	500

The VSM_{DLL} is higher than the VSM_{FN} , which can be explained by that the VSM_{DLL} is more stable close to the system collapse point. The computation of the VSM_{FN} stops at roughly the same level of system stress for every case, caused by the Newton-Raphson algorithm not converging. The VSM_{DLL} , which is essentially a more stable method to compute the VSM than the VSM_{FN} , is capable of stressing the system closer to limit before any of the stability criteria is met.

The differences between the DVSM and the VSM_{DLL} are generally smaller for Case B when the three-phased fault is cleared in 0.04 seconds rather than 0.1 seconds. This result can be referred to the previously developed theory of the transient P-V curves. The lowered system voltages during the fault-time will affect the transient P-V curves in the same manner that the tripped line did in the example developed in Section 5.1.4, but to an even larger extent. During the fault-time, the voltages drop significantly more than if the line would have simply been tripped. Thus, if the fault-clearing time is sufficiently long, it is possible that the load restoration in the system has time to push the system out of stability before the voltage restoration devices have time to act.

Table 5.2: Case B: Three-phased fault during 0.1 seconds cleared by tripping line.

Case number	<i>Constant</i>			VSM_{FN} [MW]	VSM_{DLL} [MW]	DVSM [MW]
	MVA (P/Q) [%]	I (P/Q) [%]	Z (P/Q) [%]			
1A	100/0	0/0	0/100	298	390	140
2A	99/0	1/0	0/100	298	390	190
3A	95/0	5/0	0/100	298	390	370
4A	90/0	10/0	0/100	299	400	370
5A	80/0	20/0	0/100	298	405	410
6A	50/0	50/0	0/100	299	430	390
7A	100/0	0/50	0/50	299	380	110
8A	95/0	5/50	0/50	299	385	230
9A	90/0	10/50	0/50	299	380	330
10A	50/50	50/50	0/50	299	425	330
11A	0/0	100/0	0/100	299	425	460
12A	0/0	50/0	50/100	299	455	490
13A	33.33/0	33.33/0	33.33/100	299	450	450
14A	0/0	20/0	80/100	299	490	510
15A	20/0	0/0	80/100	299	475	500

CHAPTER 6

Fast Dynamic Voltage Security Margin Estimation

This chapter presents a method for fast DVSM estimation, based on the methodology and the results established in Paper I. The method is proposed to be included as a preventive monitoring application in the developed RVSAT, earlier presented in Chapter 4.

6.1 Introduction

In Chapter 5, the circumstances when the DVSM is to prefer to the conventional VSM was presented. However, DVSM estimation is computationally demanding, where multiple time-domain simulations are required to trace the security limit for a range of different contingencies. In this chapter, a methodology for fast estimation of the DVSM is proposed to overcome the computational difficulties when estimating the margin. The method use NNs to provide both an estimate of the actual DVSM at a specific OC, and to determine the dimensioning contingency for the system with respect to the DVSM. These *estimated* values are then used as starting points in a method called dual binary search to significantly reduce the required computational time in computing the *actual* DVSM. The method is developed to mitigate inconsistency issues associated with ML methods under new or unseen operating conditions.

6.2 Methodology for fast estimation of the DVSM

The NNs are trained on a training set consisting of i) computed values of DVSM for a range of different OC, and ii) the dimensioning contingency for the same OCs. Once the NNs are trained, they can almost instantaneously provide estimations of the DVSM for a certain OC, and classify which contingency is most likely to

be dimensioning for the system. The first step of the method is the generation of credible operating conditions (OCs) and estimations of the DVSM for a set of credible contingencies. The method is tested on the Nordic32 test system found in Fig. 4.2 with all data and models as presented in [65]. After a representative training set has been generated, the training scheme of the two NNs is presented. Each step in the methodology is described in the following subsections.

6.2.1 Generation of training data

The training data for the NNs were generated using PSS[®]E 34.2.0 with its in-built dynamic models [69]. Here, full time-domain simulations have been used, but the methodology could also be generalized for combinations of QSS and full time-domain simulations. The steps of generating the training data are illustrated in the flowchart in Fig. 6.1 and can be summarized as follows:




Figure 6.1: Flowchart of the generation of training data for the DVSM and the neural network.

- (i) **Choose initial operating conditions:** All initial OCs were randomly generated around the stable operating point of the simulated Nordic32 system denoted as "operating point B" in [65]. The total load in the system for each initial OC was generated by multiplying *all* the active loads randomly from the *same* uniform distribution (80 % of the original load as the lower limit, 95 % of the original load as the upper limit). Then, each *individual* load was varied by again multiplying the now updated load value with a random variable

generated from a new uniform distribution (this time with 80% as the lower limit, and 120 % as the upper limit). The power factor of all loads were kept constant. The total change in loading was then randomly distributed among all the generators in the system. The generated initial OCs were first solved using a conventional full Newton-Raphson load flow solution, which served as a starting point for the dynamical simulation. In the case the system was not found feasible, the initial OC was re-initialized.

- (ii) **Increase system stress and solve load flow:** The system stress was then increased for the secure initial OC by increasing the power transfer between the two areas "North" and "Central". The increased system stress was achieved by increasing the loads in the "Central" area with a total of $\Delta P_1 = 200$ MW, while simultaneously increasing the generation in the "North" area with the same amount. The power factors of each load were again kept at the initial values. The distribution of the added load and generation was based on the initial load or the rated capacity of each generator. Thus, a bus with a larger initial load, or a generator with a higher rated capacity, received a larger share of the increased load and generation. All generation that could not be supplied by the regular generators were distributed to the slack bus generator in the system, g20, see Fig. 4.2. After the loads and generation were updated, the load flow was reiterated which then served as a starting point for the time-domain simulations. To avoid numerical and stability issues when increasing the system stress of the static system, the system stress was increased in small increments where a load flow solution was solved for each increment.
- (iii) **Run time-domain simulation and test for security:** A time-domain simulation was then initialized for the first contingency. In the relatively small Nordic32 test system, the same single contingency was found to be dimensioning for almost all different initial OCs. To test the possibilities of using a NN to classify the dimensioning contingency, two different contingencies were hand-picked as they were found to be dimensioning for different OCs. The tested contingency type was a three-phased fault on a transmission line during 0.1 seconds, followed by tripping the faulted line which was then kept tripped during the remaining time of the simulation. The lines between the buses "4031 - 4041" and "4032 - 4044", connecting the "North" and "Central" areas were used, see Fig. 4.2 for reference. Each simulation then ran for a total of 500 seconds. The system was considered secure if, at the end of each simulation, *all* transmission bus voltages were above 0.90 pu.

Each dynamic simulation ran for a total of 500 seconds but was in the case of a major voltage collapse stopped in advance. The simulation time was chosen to ensure that the system either fully stabilized or collapsed. It should be noted that the required simulation time is dependent on the power system in consideration, and it is likely that different simulation times would be required in actual implementations of the algorithm. The system was considered secure if, at the end of each simulation, *all* transmission bus voltages were above 0.90 pu.

- (iv) **Re-iterate and test other contingencies:** In case the system was found secure for the tested contingency, the system stress was increased again with ΔP_1 , followed by another security test. In case the system was *not* found secure, the previously added system stress was halved, and the process was re-initialized. This process of iteratively updating the system stress and testing for security continued until the increase in system stress was below a precision value of $\epsilon = 5$ MW. The DVSM was then computed by taking the difference in system loading between the initial OC and the secure system with the highest level of system stress.

Once the DVSM for the first contingency was computed, the same procedure was repeated for the second contingency. To save computational time, the estimated DVSM for the *first* contingency was used as a *starting point* for the estimation of the second contingency. If the system at that level of system stress was found secure for the second contingency, the simulation was stopped. Otherwise, the search algorithm continued until a new smaller value of the DVSM was found.

- (v) **Sampling the input values and target values:** An input vector \mathbf{x} consisting of measurements of all bus voltage magnitudes and angles, and active and reactive power flows were sampled from each one of the initial OCs. The choice of which input values to include in the training was based on the results in [70], which found that bus voltage magnitudes and angles were found to be the best combination of inputs when estimating the VSM using a NN. The active and reactive power flows were then added as additional inputs as this was found to increase the accuracy in the estimations even further. Two target vectors \mathbf{y}_{DVSM} and \mathbf{y}_{Cont} were generated by sampling the DVSM for each case, and the contingency that was dimensioning for the specific case, respectively. The previously described steps were re-iterated until a sufficiently large training set was generated. Due to the random nature in which the training data was initialized and generated, some of the generated OCs were found to be correlated with very low DVSM values, despite being initialized with low system loading. To ensure that no anomalies were included in the training set, all OCs resulting in DVSM values below 150 MW was excluded from the training set.

6.2.2 Architecture of the neural networks

The architecture of the two NNs used in this thesis are presented in Fig. 6.2 and the specific details regarding the architecture and the training parameters of the two NNs are specified in Table 6.1. In the training phase, the two NNs takes the same vector of input values, which are forwarded to each of the hidden layers through a set of weights, illustrated by the lines connecting each of the neurons. The output of each neuron in the hidden layer is computed using a non-linear activation function on the sum of all the inputs, which is then forwarded to the output layers. The rectified linear activation function (ReLU) was used as the non-linear activation function for

the two NNs. For the NN estimating the DVSM, the outputs are forwarded to a regression layer with a linear activation function. For the NN responsible of ranking the contingencies, the outputs are forwarded to a layer with a softmax activation used for classification. The softmax activation function is generally used for multi-class classification but generally works well also for binary classification as is the case here. The softmax activation function outputs a probability vector, where each class is given a certain probability. The probability vector can then be used to rank the contingencies in order of which most likely will become dimensioning.

In the training phase, the networks use the true target vectors y_{DVSM} and y_{Cont} , while during the test or prediction phase, the network estimates the DVSM and the ranked contingencies by generating the vectors \hat{y}_{DVSM} and \hat{y}_{Cont} for the current OC. The supervised training approach aims to update and learn the suitable values for the weights connecting each layer, implicitly modeling the non-linear relationship between the inputs and outputs.

6.2.3 Training

Different data sets were used in training, validation, and testing of the method. The training data has the dimension (364 x 6000), where the dimension represents the number of inputs, and the total number of training cases, respectively. Each network was trained for a maximum number of epochs, where an epoch is finished when all the cases in the training set have been used to update the network parameters. To reduce overfitting on the data, ridge regression (also known as L2 regularisation) was used to ensure the data does not rely too heavily on any single feature. To further reduce overfitting, a technique called dropout was applied where a certain percentage of the connections between each layer were masked/dropped, to ensure that the model does not rely too heavily on certain connections. The mean squared error (MSE) was used as a metric for the NN estimating the DVSM, while the categorical cross-entropy loss function was used for the NN classifying the dimensioning contingency. An adaptable algorithm for gradient-based optimization, Adam, was used in training the network [71]. The learning rate was the only parameter that was specifically tuned for the algorithm, while the remaining used the default values according to [71].

It should be noted that both the training parameters and the architecture of the two networks have been iteratively tuned to increase the regression and classification accuracy. A deeper architecture with more hidden layers was found to not increase the performance for the specific test case and training set size. Other hyperparameters and network architectures would likely have better performance for other test systems than the Nordic32. By increasing the training set size further and spending even more effort in tuning the networks, an even better accuracy could be achieved.




Figure 6.2: The two ANNs trained to evaluate the DVSM and the buses with the lowest margin to instability.

6.2.4 Fast DVSM estimation and dual binary search

In [11], a binary search method was proposed to estimate the DVSM. Here, an alternative approach denoted as the *dual binary search method* is proposed, which should be able to increase the computational speed of the DVSM. The trained NNs in Fig. 6.2 take the same set of measurements and generate: (i) an estimated value of the DVSM, and (ii) an estimated ranked order of the contingencies that most probably will be dimensioning for the current OC. The estimated DVSM is used as a qualified estimate of the real DVSM, which is validated through actual time-domain simulations. The dual binary search method is then used to take advantage of the estimated DVSM and the dimensioning contingency to reduce the computational time when validating the real DVSM for the system.

The dual binary search method is illustrated for two cases in Fig. 6.3. Case 1 illustrates the estimation process for an overestimated value of the DVSM, while Case 2 illustrates the estimation process for an underestimated value of the DVSM. Black dots indicate secure operating points and white dots indicate insecure operating points. The estimated DVSM is always the starting point for the search of the *ac-*

Table 6.1: Design and hyperparameters used in training.

	Parameter	NN 1 / NN 2
<i>Data</i>	Number of inputs	364 / 364
	Training cases	6 000 / 6 000
	Validation cases	400 / 400
	Test cases	400 / 400
<i>Architecture</i>	Hidden layers	1 / 1
	Final activation function	Linear / Softmax
	Hidden cells	128 / 32
	Hidden layer activation	ReLU / ReLU
<i>Training</i>	Max Epochs	1 000 / 3 000
	Learning rate (α)	$1 \cdot 10^{-6}$ / $1 \cdot 10^{-5}$
	Dropout	0 % / 50 %
	L2 parameter	0.01 / 0.01
	Optimizer	Adam / Adam [71]
	Loss metric	MSE / Categorical cross-entropy

tual DVSM of the system. The system stress is increased to this point iteratively using the approach explained in section (ii) to avoid convergence problems.

Once the stressed static base case is found, a time-domain simulation is initiated for the *highest* ranked contingency by the second NN, which is the contingency that most likely will be dimensioning for the DVSM. The initial estimated DVSM level is then tested for the chosen contingency. In case it is stable (respectively unstable), the system stress is increased (respectively decreased) with a certain value represented by ΔP_2 . A value of ΔP_2 equal to the mean squared error (MSE) of the estimated values for the DVSM is proposed, which should represent a reasonable uncertainty and step size for the estimation. If the new operating point is found to be secure, the system stress is again increased with ΔP_2 . In case it is not found to be secure, which happens in the example illustrated in Fig. 6.3, the system stress is reduced by $\Delta P_2/2$. The dual binary search is then continued until a secure operating point is found and when the step size in system stress change is smaller than a specified precision level (ϵ).

This level of system stress is then tested for the other contingencies, in ranked order, until all lower-ranked contingencies have been tested and found secure. For both the cases illustrated in Fig. 6.3, this level of system stress for the second-ranked contingency was found to result in a secure operating point. A third and final ranked contingency is then tested, which for Case 1 in Fig. 6.3 is found to be insecure. The system stress is thus reduced further for this case, resulting in a secure operating point which then constitutes the dimensioning DVSM for the system. It should be noted that for Case 1, the contingency ranking was not perfect, with the result that an extra time-domain simulation was required.

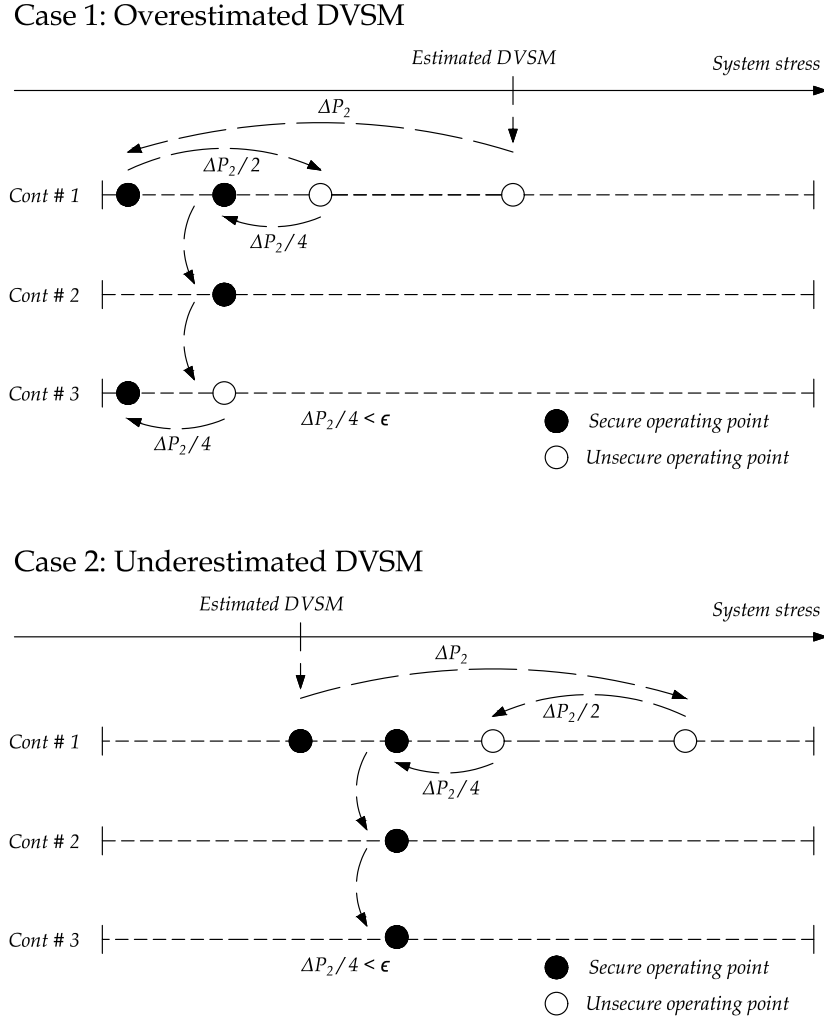


Figure 6.3: Dual binary search for multiple contingencies.

6.3 Results and discussions

In the following section, the results from the regression and classification on the test set for the two NNs are presented. Furthermore, the reduction in computational effort is compared between the conventional tracing method and the proposed dual binary search method. Finally, practical applications and discussions related to DVSM estimation are presented.

6.3.1 Regression and classification accuracy

The prediction accuracy for the NN estimating the DVSM is presented in Fig. 6.4, where the estimated DVSM is plotted with respect to the real DVSM for the test set. The diagonal line indicates where the points should lie in case the estimated DVSM perfectly matches the real DVSM. Table 6.2 lists the mean and maximum




Figure 6.4: Prediction accuracy of estimating the DVSM.

Table 6.2: Regression results for estimating the DVSM.

Mean estimation error	Maximum estimation error	Mean squared error
1.49 %	10.96	13.35 MW

error of the estimations in percentage, as well as the MSE presented in MW. The results indicate that the NN is generally capable to accurately estimate the DVSM given an initial OC, with a mean error for the test set of 1.49 %. The maximum estimation error was found to be 10.96 %, while the MSE was estimated to 13.35 MW.

The classification accuracy of the NN used in ranking the dimensioning contingency is presented in Table 6.3 in the form of a confusion table. Each number in each row represents the instances of the real dimensioning contingencies, while each number in each column represents the instances of the predicted dimensioning contingencies. The conditional probabilities of correctly classifying the dimensioning contingency are presented in the column furthest to the right. Similarly, the conditional probabilities of a dimensioning contingency actually *belonging* to the predicted class are presented in the bottom row of the table. The total accuracy for the classification is presented in the rightmost corner of the table, and an accuracy of 91.3 % was provided for the test set. Thus, in about 9 instances of 10, the NN is capable of classifying which contingency that will be dimensioning for the DVSM for a specific OC. It should again be mentioned that the estimation and classification results could be enhanced further by either increasing the training set size, or by a more careful exploration of suitable hyperparameters for the training of the networks.

Table 6.3: Confusion table showing real and predicted dimensioning contingencies.

		<i>Predicted</i>		Accuracy
		L4044 - L4032	L4044 - L4031	
<i>Real</i>	L4044 - L4032	104	23	81.9 %
	L4044 - L4031	12	261	95.6 %
	Accuracy	89.6 %	91.9 %	91.3 %

Table 6.4: Reduction in computational effort using the proposed method.

<i>Average number of time-domain simulations</i>		
Conventional tracing method	Dual binary search method	Relative reduction in computations
15.3	4.7	-69.2 %

6.3.2 Computational efficiency

In this section, the computational efficiency is compared between the proposed fast dual binary search method and the more conventional tracing method that was used in generating the training set (see Section 6.2.1 for reference). The proposed fast dual binary search method, explained in Section 6.2.4, uses the estimated DVSM value and the dimensioning contingency as a starting point to validate the real DVSM. The computational efficiency is measured as the average number of time-domain simulations required in estimating the DVSM. The results of using the two different methods are presented in Table 6.4. The average number of time-domain simulations required in estimating the DVSM using the *conventional* tracing method was found to be 15.3, while the corresponding number using the proposed dual binary search method was 4.7. The reduction in average number of time-domain simulations required was thus -69.2 % when the proposed method was applied.

It should be noted that the *exact* comparison in computational efficiency between the two methods is of comparatively little interest, as it mainly applies to the specific test case used in this thesis. For instance, the computational savings are probably significantly higher in most real applications, where a larger range of contingencies may be dimensioning for the DVSM. Furthermore, in real applications where the range of the DVSM may be larger than what has been used here, the conventional search algorithm would require significantly more time-domain simulations to find the true DVSM. Similarly, it is also possible to further enhance both the conventional search algorithm and the dual binary search algorithm by, for instance, choosing more suitable values of ΔP_1 , or increasing the precision value of ϵ . The most notable result is instead that the computational effort in estimating the DVSM can be reduced from requiring a large number of time-domain simulations, to only

requiring a few. Although a few time-domain simulations would still take some time to compute for a real power system, it should be possible to provide sufficiently fast estimations of the DVSM to classify it as a "*near real-time*" estimation.

6.3.3 Impact of sudden topology change

In any real application, the performance of a NN is dependent on its generalization capability. This refers to the capability of the NN to generalize the learning from the actual training set to other, yet unseen, cases. In this section, the performance of the NNs to generalize their estimations when subjected to test cases where unplanned topology changes have taken place is examined.

For simplicity, only topology changes in the form of opened transmission lines are considered. To ensure that the Nordic32 test system is still secure despite the topology changes, only topology changes in the "North" region, see Fig. 4.2, were considered. Furthermore, only transmission lines between buses served by two parallel transmission lines were used in generating the test set with topology changes. A new test set of 400 cases was then generated in the same manner as explained in Section 6.2.1, with the difference that the topology changes were now added randomly.

In Fig. 6.5, the prediction accuracy of the NN estimating the DVSM is presented when the network has been trained on two different training sets. Case 1 presents the prediction accuracy when *no* unseen cases with topology change have been included in the training set. Case 2 presents the prediction accuracy when a few (100) training cases with topology change have been included in the training set. For Case 2, the two NNs were re-trained on the updated training set using the same training approach as previously described. Table 6.5 lists the mean error of the estimations in percentage for each case, as well as the MSE presented in MW. The result for Case 1 indicates that a sudden topology change will significantly affect the accuracy of the predictions. Although many cases were accurately predicted, the number of outliers increased significantly. The prediction accuracy was higher for Case 2, even though only a very small number of cases with the topology changes were added to the training set. The classification accuracy of the NN used in ranking the dimensioning contingency was also affected significantly for the two cases, with a total classification accuracy of 55 % and 78.5 % for case 1 and case 2, respectively.

The results highlight the importance of obtaining a representative training set and also taking into account the possibilities of unplanned topology changes. It was seen, that by the inclusion of even a very small set of training cases with various topology changes, the prediction accuracy could be increased significantly. Thus, in the event of an unplanned change in the system, the system operator could quickly generate a small training set on the new OC, and then retrain the NNs on the generated data. It should be noted that the proposed robust methodology of always validating the estimations of the NNs with *actual* time-domain simulations reduces the impact of these types of erroneous estimations. The main impact of a poor estimation of the DVSM will be that the time to validate it will increase.

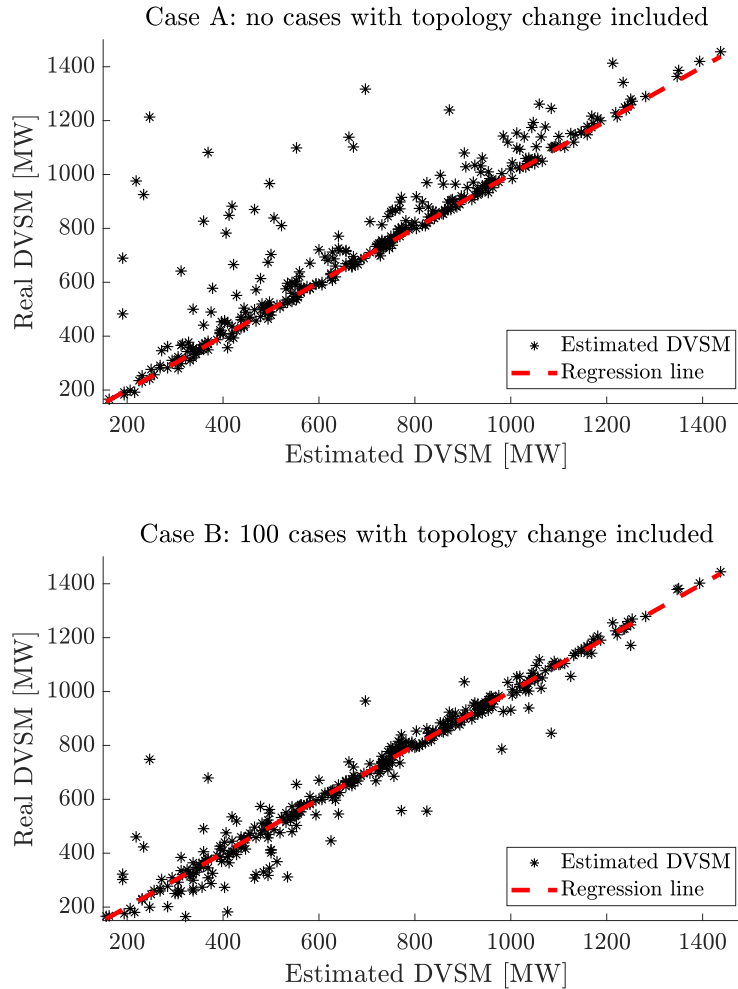


Figure 6.5: Prediction accuracy of the NN during sudden topology changes. With and without trained on a training set with a 100 training cases with topology changes included.

6.3.4 Discussions and practical applications

The proposed method is aimed to be used as an online tool for system operators to estimate a power system’s dynamic voltage security margin. The method does not necessarily have to replace conventional VSM estimation, but may instead be used as an additional source of information to system operators to provide better and more accurate estimates of the total transmission capacity in their systems.

Theoretically, the DVSM estimates by the NNs could be used directly to provide real-time estimates of the security margin. However, despite years of research, examples where these methods have been practically applied in system operators’ monitoring and control systems are, to the authors’ best knowledge, very few. From a system operator view, an inferior method that *always* works are generally preferable to a

Table 6.5: Regression results of the DVSM estimation for the two cases.

	Mean estimation error	Mean squared error
Case 1	12.97 %	134.38 MW
Case 2	7.33 %	63.59 MW

superior method that in some instances does not. The proposed method is thus suggested to utilize the advantages of ML, while still ensuring that the method always provides good estimates regardless of the current OCs.

The results in the previous section indicated that by using the proposed method, the number of required time-domain simulations to estimate the DVSM could be reduced to only a few, allowing system operators to estimate the DVSM in a time frame that could be defined as "near real-time". The update frequency of security margins will affect the required transmission reliability margins as the underlying system continuously change between the assessments. The actual estimation speed will still be affected by a range of different factors such as the computational speed of the hardware being used, the size of the specific power system in consideration, or the required precision (i.e. the value of ϵ).

Measurements of bus voltage magnitudes and angles, as well as active and reactive power, have been assumed to be available, either directly from measurements or from state estimations of the system. However, to ensure that missing values and errors are filtered out, all measurements should preferably be preceded by a state estimator. To adapt to the evolving operating conditions and self-rectify any bad predictions, the two NNs should be trained continuously during operation. Using approaches such as stochastic gradient descent, the NNs weight parameters can continuously be tuned to increase the robustness and accuracy.

CHAPTER 7

Voltage instability prediction using a long-short term memory network

This chapter presents an algorithm for real-time VIP, based on the developed methodology and the results established in paper II and paper III. The methodology for training the VIP tool to indicate where in the system the instability would emerge was first presented in paper III. The final VIP tool is mainly based on the proposed architecture presented in paper II. The method is proposed to be included as an emergency monitoring application in the developed RVSAT, earlier presented in Chapter 4.

7.1 Introduction

In Chapter II, the voltage instability event and the need for emergency monitoring systems were discussed. It was further discussed that conventional methods for VID may be too slow to detect instability in time for system operators to initiate sufficient emergency control actions. An alternative option is to use ML-based methods for VIP, which can predict the onset of instability only seconds after a disturbance has occurred in a system.

These methods are trained to indirectly correlate a post-disturbance state and learn its dynamical trajectory, to be able to directly assess whether the system is heading towards instability. Most previously developed methods for VIP have in common that only instantaneous measurements are used as inputs to the VIP algorithms. These inputs represent the "state signal" that the ML algorithm uses to predict the future state. Ideally, the state signal should summarize all relevant information required to determine the future state of the system. A state signal achieving this is said to have Markov property [72]. However, the dynamic response of a power system cannot be modeled as a first-order Markov process using only the static states provided by available measurements in the power system. Rather, the future

state of the system also depends on a range of unknown state variables such as the rotor speed of generators, tap positions, or rotor slips of induction motors.

In response to these limitations, a new method based on a recurrent neural network (RNN) with long short-term memory (LSTM) is proposed. LSTM networks excel at capturing long-term dependencies [64], which is an inherent aspect in long-term voltage stability [10]. The methodology and test results for the proposed method, from here on denoted as LSTM-VIP, are presented in the following sections.

7.2 Methodology for LSTM-VIP

The proposed method for real-time VIP is based on off-line training of an LSTM network on a large data set consisting of time-domain simulation responses following a set of credible contingencies. The method is aimed to be used as a supplementary warning system that can assess the current state of the system in real-time. The LSTM network takes real-time and historic measurements and attempts to assess whether the *current* state will cause voltage stability issues several minutes into the future. As time progresses and if new events occur in the system, the network updates the assessment continuously. The network is also adapted to be able to indicate *where* in the system instability emerges, following the approach developed in [73], allowing more cost-effective countermeasures.

The first step of the method is the off-line generation of credible operating conditions (OCs) and contingency scenarios using time-domain simulations. The method is generic, but is here tested on the Nordic32 test system with all data and models as presented in [65]. After a representative training set is generated, training of the LSTM network is performed. Each step in the methodology is described in the following subsections.

7.2.1 Generation of training data

The generation of a training set is a critical step and a range of different initial OCs and contingencies were included to generate a representative training set. Dynamic simulations were performed using PSS®E 34.2.0 with its built-in models [69]. The steps of generating the training data are illustrated as a flowchart in Fig. 7.1 and can be summarized as follows:

- (i) **Initial OCs:** For the Nordic32 system, the initial OCs were randomly generated around the stable operation point denoted as "operating point B" in [65]. A large number of possible OCs were simulated by randomly initiating the loads from a uniform distribution around the base case load levels (80 % of the original load as a lower limit and 120 % as an upper limit), while the power factor of the loads was kept constant. The total load change was distributed among the generators based on a weighted random distribution, where a higher rated capacity of a generator results in a higher probability to cover a larger

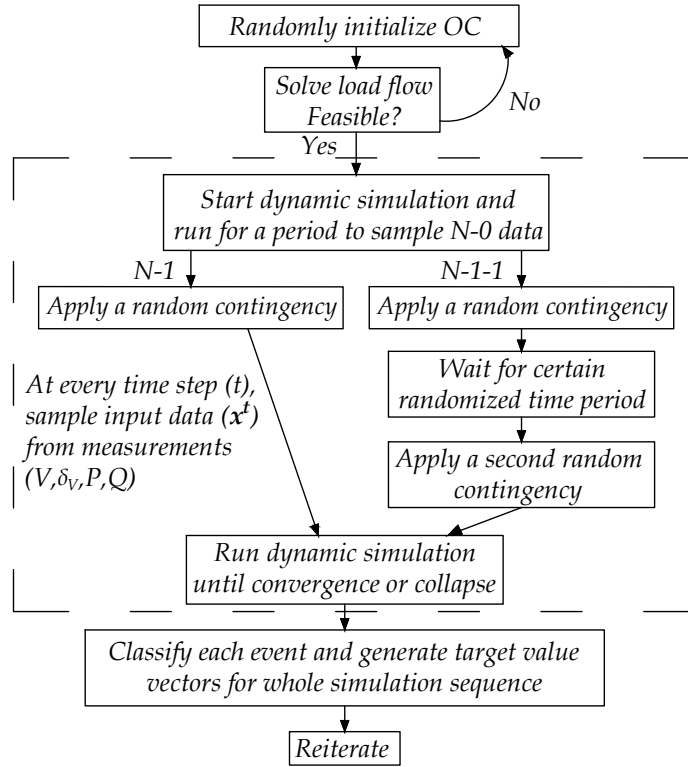


Figure 7.1: Flowchart for generating input data and target values.

share of the total load change. All generation that could not be supplied by the regular generators were distributed to the slack bus at g20, see Fig. 4.2.

In real applications, more delicate methods for efficient database generation and more careful generation of relevant OCs should be used [15, 74], where for instance the impact of unit commitment and topology changes are taken into account.

- (ii) **Solve and check for feasibility:** The generated OCs were solved with a power flow simulator, which served as a starting point for the dynamical simulation. If the system load flow did not converge, the initial OC was re-initialized.
- (iii) **Start dynamic simulation and introduce contingencies:** Two separate dynamic simulations were then initiated for the $N-1$ and the $N-1-1$ cases. The process is illustrated in Fig. 7.2. For each of the two cases, the system runs without any contingencies for 65 seconds to generate a sufficient amount of $N-0$ data for the LSTM network to train on. At $t = 66$ seconds, the *same* first contingency was applied to both of the cases. After an additional uniformly distributed random time in $[10 - 30]$ seconds after the first contingency, a secondary consecutive contingency was applied for the $N-1-1$ cases. Events resulting in several (near-)simultaneous contingencies were not taken into account ($N-k$ events).

The considered contingencies in the simulations were either (i) tripping of a generator, or (ii) a three-phased fault during 0.1 seconds, followed by tripping the faulted line, which was then kept tripped during the remaining time of the simulation. The first contingency was chosen to be a major fault, meaning a fault on any transmission line connecting the different main areas in the system (excluding the "Eq." area, see Fig. 4.2), *or* any larger thermal generator in the "Central" area. The second contingency, for the $N-1-1$ cases, included tripping of *any* transmission line in the whole system, excluding lines in the "Eq." area. No variations of load and generation were taken into account during the dynamic simulations as these, in the relatively short time of the simulations, are presumed to have a small impact on the system stability. In real settings, depending on the system and the experience of the operator, all relevant contingencies can be used in the training.

- (iv) **Sample inputs and run until stopping criteria:** For each of the two cases, an input vector \mathbf{x}^t consisting of measurements of all bus voltage magnitudes (V_{mag}) and angles (V_θ), active and reactive power flows (P_{flow} , Q_{flow}), were sampled every second ($\Delta t = 1s$) and saved in a data file. The value of $\Delta t = 1s$ is dependent on the possible measurement update rate in the actual system and will determine the rate the estimations are available to the system operator. No information regarding the type and location of applied the contingencies were sampled, as this information implicitly can be learned by the LSTM network. For instance, the LSTM network should be able to correlate a zero power flow in a transmission line with that line being out of service.

Each dynamic simulation ran for a total of 560 seconds, but was, in the case of a major voltage collapse, stopped in advance. The simulation interval of 560 seconds was chosen to allow time for *all* dynamic events to occur and for the system to either fully stabilize or collapse.

- (v) **Classification:** For each case, a sequence of true target value vectors $\mathbf{y}^1, \dots, \mathbf{y}^{560}$ was generated for every time step in the time-domain simulation. Each \mathbf{y}^t in these sequences represents the classification of the system if the system is allowed to run from time t up until 560 seconds without any changes to the current system. As time progresses and new events occur, the class of \mathbf{y}^t may change. The sequences consist of multidimensional vectors where the actual class is encoded using one-hot (binary) encoding.

The classification was performed according both to the severity and the location of the system degradation at the *end* of the time-domain simulation. The system was defined as stable if *all* transmission bus voltage magnitudes were above or equal to 1 pu, in an alert state if *any* transmission bus voltage magnitude ranged between $0.9 < V < 1.0$ pu, and in an emergency state if *any* transmission bus voltage magnitude was below 0.9 pu. Overvoltages were not taken into account.

The target values for the *alert* cases were also classified according to *where* the lowest bus voltage magnitudes were found at the end of each dynamic




Figure 7.2: Example of classification of an $N-1$ and an $N-1-1$ case.

simulation. The Nordic32 test system is predefined into four different regions, namely: "Eq", "North", "Central", and "South" [65]. The regions "North", "South", and "Eq." were found to be stable regions, and no alert events were found in these regions for any of the simulated cases. To test the capability of the network to also indicate where instability emerges in the system, the "Central" area was divided into three separate regions (indicated by **C1**, **C2**, **C3** in Fig. 7.3). The classification for each time step of each simulation belonged then to one of 5 different possibilities. Either the whole system was predicted stable; it ended up in an emergency state; or an alert state was predicted in one of the three defined regions (**C1**, **C2**, or **C3**) where the *lowest* occurring transmission bus voltage was found.

The classification process is illustrated in Fig. 7.2. The target values are always classified as stable up until the first contingency. From different combinations of OCs and contingencies, the system may then end up being in a stable state, an alert state in area **C1**, **C2**, or **C3**, or in an emergency state. For the $N-1$ case, the sequence of true target value vectors from the time of the contingency to the end of the simulation are classified depending on which of these five states the systems end up in. For the example of the $N-1$ case in Fig. 7.2, the system ends up in an alert state in the **C1** area. For the $N-1-1$ case, the target values are classified as stable up until the first contingency. The target values are then gathered from the $N-1$ case, using the end state of that simulation for classifying the state *between* the first and the consecutive contingency. After the second consecutive contingency, the system runs until it either collapses or until 560 seconds. Depending on this final state, the




Figure 7.3: One-line diagram of Nordic32 system with subareas.

sequence of true target value vectors from the second contingency until the end of the simulation is classified. In the example in Fig. 7.2, an emergency state is reached. Note that the scales in Fig. 7.2 are different from those in the simulations for easier interpretation.

It should be noted that the classification of the different states (stable, alert, emergency) could be performed more intricately to satisfy other criteria of stability. For instance, these could be related to a minimum level of loadability of the system in its post-disturbance state. The loadability limit could then be computed by, for instance, parameterized continuation methods such as the continuation power flow (CPF) method [12], or by certain line indicators [39]. Other stability criteria could include the capability of the system in its post-disturbance state to handle yet another disturbance.

- (vi) **Reiteration:** The described steps are reiterated until a sufficiently large training set is generated.

7.2.2 Architecture of the LSTM network

The proposed LSTM network architecture, shown in Fig. 7.4, is generally referred to as a "many-to-one" architecture, where previous measurements in the time sequence are used for the classification in the final block. The network consists of three stacked




Figure 7.4: The proposed LSTM network architecture.

LSTM layers which are used to capture different levels of features from the inputs. Each LSTM block consists of 32 individual LSTM cells. The first layer of LSTM-blocks takes a generated sequence of input vectors as inputs; then by mathematical operation as presented in Section 3.3.1, the output of each block is forwarded both to the following block in the sequence, as well as to the upper layer of LSTM-blocks. The inputs to the deeper layers consist only of the hidden states of LSTM blocks of previous layers, while both the hidden state and the cell state memory is passed along the time sequence between LSTM blocks of the same layer.

The LSTM network is designed to take sequences of 60 time steps of measurements as inputs. The internal architecture of each LSTM cell and functionality of the nonlinear gating units as presented in Section 3.3.1, allows the LSTM network to fully utilize and pass forward the information from the first to the final time step in the sequence. The third layer of LSTM-blocks only passes the output forward along the time sequence. The output layer at time t is a fully connected network with softmax activation for classification. In training, the network uses the true target vector \mathbf{y}^t at time t , while during the test or prediction phase, the network estimates a prediction vector $\hat{\mathbf{y}}^t$ at time t . The interpretation of the prediction problem is further explained in section 7.2.4.

7.2.3 Training the LSTM network

Different data sets were used for training, validation, and testing of the method on a mix of $N-1$ and $N-1-1$ cases. The training data set has the following dimension ($135,000 \times 364 \times 560$), where the dimension represents the number of training cases, the number of inputs, and the total interval in seconds for each simulation, respectively.

Before training, a process generally referred to as sequence preprocessing was per-

Table 7.1: Design and hyperparameters used in training.

	Parameter	Values and size
<i>Data</i>	Simulation interval	560
	Input dimension	364
	Input data type	$V_{mag}, V_{\theta}, P_{flow}, Q_{flow}$
	Target classes	5
	Training cases ($N-1+N-1-1$)	45,000 + 90,000
	Validation cases ($N-1/N-1-1$)	5,000 / 10,000
	Test cases ($N-1/N-1-1$)	10,000 / 10,000
<i>Architecture</i>	LSTM layers	3
	LSTM sequence length	60
	FC activation function	Softmax
	LSTM hidden cells	32
	LSTM Activation function	Tanh
<i>Training</i>	Max Epochs	500
	Learning rate (α)	0.0001
	Dropout & recurrent dropout	50 % / 50 %
	Optimizer	Adam [71]
	Loss metric	Categorical cross-entropy

formed to prepare batches of sequences with suitable length. The network is designed to take a sequence of 60 time steps of measurements as inputs and subsequences with a length of 60 time steps ($\mathbf{x}^{t-59}, \dots, \mathbf{x}^t$) were thus extracted from the 560 seconds long simulation intervals, for different values of t . For each subsequence of input vectors, a corresponding target value (\mathbf{y}^t) at time t was gathered. The sequence pre-processing was performed 120 times for *each* training and validation case by varying t between values of $t = [60, 180]$. The lower bound of t is required to always allow historic data to be included into the sequence. The LSTM network could have been trained on the whole simulation interval by increasing the upper bound of t from 180 to 560. However, since the method is proposed to be used in fast VIP applications, there is less usefulness of predicting instability long after the contingencies have occurred.

The generated subsequences were then used to train the LSTM network. Due to memory limitations, a method called mini-batch gradient descent was utilized where mini-batches of 1000 subsequences were used separately to train the network. The training was performed for a maximum of 500 epochs. An epoch is finished when all generated batches have been used to update the network parameters. An adaptable algorithm suitable for gradient-based optimization of stochastic objective functions, more commonly known as "Adam" was used in training the network [71]. The algorithm used default parameters according to [71], except for the learning rate which was tuned. The loss function on which the optimizer is applied is the categorical cross-entropy function, which is suitable for multi-classification problems. To avoid overfitting the data, two regularization techniques were used during the training. First, early stopping was implemented, and the training of the network was stopped

in case the performance on the validation set did not improve after six epochs. Second, a technique called dropout was applied, where a certain percentage of the connections between inputs and the LSTM cells were randomly masked (or "dropped") to reduce overfitting on the data. Both conventional dropout and recurrent dropout between consecutive blocks were applied during the training phase.

All other parameters related to the training of the network are presented in Table 7.1. The LSTM network was trained and implemented in Python, using the Keras library with TensorFlow backend. The architecture and parameters used to train the network have been iteratively tuned to increase the classification accuracy. It should be noted that the tuning could be extended even further to allow an even better classification accuracy.

7.2.4 Interpretation and intuition of the VIP problem

By the proposed training and architecture of the LSTM network, a classification problem is solved where the *current* system state space is separated into different regions. Every state on a trajectory to a stable, alert (in **C1**, **C2**, or **C3**), or emergency state is labeled accordingly. The LSTM network is then trained on this data to implicitly learn these asymptotic properties of solutions and the trajectories of the system state. Once trained, the network can correlate the inputs, current and historic measurements, with a certain state-space region and trajectory, allowing warnings of voltage instability only moments after a contingency have occurred in a system. The classification is performed under the assumption that the current system is unchanged, meaning that no additional contingencies or changes in generation and load configuration will occur. However, as time progresses, new observations are used as inputs to the LSTM network to continuously update and incorporate such changes in the system.

This VIP problem should be interpreted as a fixed horizon prediction problem, where the prediction horizon always is the final state given by the trajectories of the (dynamical) system. This interpretation assumes that the simulation horizon of the generated time-domain simulations are sufficiently long so that extending the simulation horizon even further, for this particular system beyond 560 seconds, would not change the partitioning of the state space.

7.3 Results

In this section, the classification accuracy of the LSTM network is presented for two different test sets, one containing only $N-1$ cases, the other containing $N-1-1$ cases. Each test set was composed of 10,000 cases of dynamic simulations. The test results of the predictions are presented using categorical accuracy, where the *indices* of the true target values are compared to the argument maxima of the predictions. The accuracy at *each* time step is then calculated over time for each of the two test sets.




Figure 7.5: Classification accuracy over time for the proposed LSTM network.

The data were fed into the network in the form of a rolling window, with subsequences generated in the exact same manner as described in Section 7.2.3. As time t progresses, new measurements entered the network from the rightmost block in the input layer and were shifted to the left in each time increment. Since the LSTM network requires a sequence of 60 time steps of data, no predictions were made before $t = 60$. To facilitate the presentation in the following figures, a new time index T is introduced here. The relationship between the two time indices is $T = t - 60$. The LSTM network's performance for VIP is not only tested during the short JAD state but during a longer period of the dynamic trajectory the system takes following the disturbances. This is performed to test the network's capability to incorporate new observations and improve its assessment as time progresses throughout a voltage instability event. The classification accuracy is only plotted for 120 seconds after T to better visualize the changes in accuracy after the contingencies.

The classification accuracy *over time* is presented in Fig. 7.5. The classification accuracy for the $N-1$ test set dropped significantly at $T = 6$ seconds, which is the same instant that the first contingency is applied. The large drop in classification accuracy can be attributed to low bus voltages instantaneously following the first contingency, which the LSTM network has learned to correlate to a voltage instability event. The large drop in accuracy only remained for a single measurement point. After the first contingency, the classification accuracy increased and remained constant at 100% for the rest of the simulations.

The classification accuracy for the $N-1-1$ test set was identical up until the time when the consecutive contingencies were randomly applied. During this time, illustrated by the arrows in Fig. 7.5, the classification accuracy decreased slightly. Since these contingencies do not occur at the same time instant in each test case, the same instant drop in accuracy as for the $N-1$ cases was not seen. The accuracy then

gradually increased and stabilized at around 97-98 %.

The results show that the LSTM network can classify and predict future stability almost perfectly for the N-1 contingency cases and with good accuracy for the N-1-1 cases. To examine which cases were misclassified, the prediction accuracy for the two test sets, evaluated at $T = 50$ seconds, are presented in Table 7.2 in the form of a confusion table. Each number in the column in the table represents instances of the predicted classes and each number in the row represents the instances of the actual classes. The (empirical) conditional probabilities of correctly classifying a certain state is presented in the column furthest to the right. Similarly, the conditional probability of a state *actually* belonging to the predicted state is presented in the bottom row of the table. The total accuracy is presented in the lower right corner of the table. The accuracy for all N-1 cases is 100 % and no cases are falsely classified. For the N-1-1 test set, the lowest classification accuracy occurred for the alert states. After inspection of the falsely classified cases, it was found that several of these were borderline cases where the transmission bus voltage magnitude used in the classification were very close to what was used in the other classes. The highest classification accuracy occurred for the emergency cases with 99.8 %.

It should be noted that the test and training sets were weighted with more cases ending up in certain classes than others. It is thus probable that the results are slightly biased with higher accuracy for these classes, and that the classification accuracy of the other classes may be lower as an effect.

7.4 Sensitivity analysis

In this section, a sensitivity analysis of various hyperparameters and other aspects that will affect the performance of the LSTM network is presented. First, the impact of the sequence length of the LSTM network is examined and compared to that of a conventional NN. The impact of different measurement update rates and is then examined, followed by a study of the network's generalization capability.

7.4.1 Impact of sequence length

In this section, the performance of the sequence-based approach is tested and compared against a conventional feedforward NN, which only uses a single snapshot of measurements as inputs. Further, to test the impact of a *shorter* time sequence, the results of an LSTM network using a time sequence of 30 time steps, instead of 60, are presented.

To allow a fair comparison between the two approaches, the feedforward NN used in this comparison was designed to be as similar as possible to the LSTM network. Essentially, the design of the NN in the comparison is identical to the *final* time step in the LSTM network presented in Fig. 7.4, with the difference that each layer consists of a hidden layer of neurons. The designed NN has thus three hidden

Table 7.2: Confusion table showing prediction results and accuracy of the LSTM network evaluated at $T = 50$ seconds.

Classification	<i>Predicted states (N-1 / N-1-1)</i>						
	Stable state All areas	C1	C2	C3	Alert state All areas	Emergency state All areas	Accuracy
<i>Actual states</i>							
Stable state	All areas	0 / 36	0 / 11	0 / 8	0 / 8	0 / 8	100 / 94.8 %
	C1	856 / 562	0 / 3	0 / 0	0 / 5	0 / 5	100 / 98.6 %
Alert state	C2	0 / 5	1874 / 1222	0 / 0	0 / 109	0 / 109	100 / 91.1 %
	C3	0 / 0	0 / 12	0 / 208	0 / 0	0 / 0	- / 94.5 %
Emergency state	All areas	0 / 0	0 / 10	0 / 0	4504 / 6649	4504 / 6649	100 / 99.8 %
Accuracy		100 / 99.6 %	100 / 93.2 %	100 / 97.1 %	100 / 98.2 %	100 / 98.2 %	100 / 97.9 %

unit layers, each layer with 32 hidden nodes. The same FC layer with a softmax activation function was used. The training for the NN was performed identically as for the LSTM network, with the exception that instead of a sequence of input values, a single snapshot was used. The LSTM network using a shorter time sequence was trained identically to that of the longer LSTM network with the exception that a shorter sequence of 30 instead of 60 time steps was used.

In Fig. 7.6, the classification accuracy on the $N-1-1$ test set is presented for the two LSTM networks with the different time sequence length and for the conventional NN. The classification accuracy for the conventional NN was around 93 % after all the consecutive contingencies had been applied, while that of the proposed LSTM network is around 97-98 %. The results show that the performance of the LSTM network using 60 time steps in the sequence significantly exceeded that of the conventional NN, generally providing better classification accuracy over the whole time frame of the simulation cases.

The classification accuracy of the LSTM network using a shorter sequence was similar to the one using a longer sequence, with the difference of a large drop in classification accuracy occurring at around $T = 46$ seconds, see Fig. 7.6. The accuracy declined for 20 seconds and was then restored to around 97 % accuracy. A similar decline in classification accuracy, though less significant, can be noted for the LSTM network using the longer time sequence at $T = 76$. Thus, a decline in classification accuracy started exactly 60 respectively 30 seconds after the consecutive contingencies were introduced (at $T = 16$) for the two networks, corresponding to the network's respective sequence length. One explanation of these results is that the LSTM networks utilize information concerning the contingency and *pre-contingency* state to enhance the classification accuracy. When the networks starts to lose the




Figure 7.6: Impact of sequence length on classification accuracy.

information about the pre-contingency state, the chance of a misclassification increases. The results strengthen the hypothesis that a long sequence LSTM network could be used to enhance the state signal to provide better classification accuracy. Theoretically, an even longer sequence could be used to increase the accuracy even further. However, this would increase the computational cost of training, and a balance between classification accuracy and computational cost should be sought.

7.4.2 Impact of measurement update rate

The performance of the LSTM network is in this section tested for different values of the measurement update rate. The performance is compared between the previously assumed available measurement update rate of $\Delta t = 1s$ and the slower update rates of $\Delta t = 3s$ and $\Delta t = 5s$. Due to the slower update rates, the architecture and the number of LSTM blocks along the time sequence had to be reduced accordingly. The original LSTM network was designed to take subsequences of 60 time steps of measurements as inputs. Thus, for the LSTM network adapted for $\Delta t = 3s$, the number of LSTM blocks along the time sequence was reduced to a third (20 blocks along the time sequence), while the number of blocks for the LSTM network adapted for $\Delta t = 5s$ was reduced to a fifth (12 blocks along the time sequence). The LSTM networks adapted for the new measurement update rates were then trained identically to the original LSTM network, with the difference that now only every third, respectively fifth, measurement in each generated subsequence were passed on the networks.

The classification accuracy for the different values of Δt is presented in Fig. 7.7 using the same $N-1-1$ test set as in previous sections. The results show that the performance when using a measurement update rate of $\Delta t = 1s$ exceeds those using a slower update rate. The largest difference can be identified during the period when the second consecutive contingencies are applied, which indicates that a lower value of Δt is especially valuable for classification during the short time that follows a disturbance. It should be noted that due to the slower update rates of Δt , there is no dip in the classification accuracy following the first contingency.

A larger value of Δt may also increase the *time* it takes to accurately predict instability, as new measurements are being passed less frequently to the LSTM network. In Table 7.3, the average time, after a contingency, to accurately predict the future state of the system is presented for the different values of Δt . The average time is only presented for the time it takes to correctly classify the system states following the *second* consecutive contingency, since correct classification following the first contingency was almost instantaneous in all test cases. The time was computed as the averaged passed time after the second contingency, up until the time when the LSTM network could consistently and accurately predict the state of the system. For cases that took *longer* time than 100 seconds to be correctly classified, a detection time of 100 seconds was assumed to avoid skewed averaged values.

The average time to correctly predict the system state was found to be 6.6 seconds




Figure 7.7: Classification accuracy over time for different values of Δt .

Table 7.3: Average time to predict the onset of voltage instability.

	<i>Measurement update rates</i>		
	$\Delta t = 1s$	$\Delta t = 3s$	$\Delta t = 5s$
Average prediction time [s]	6.6	8.7	10.7

for the proposed LSTM architecture using a measurement update rate of $\Delta t = 1s$. The corresponding values for the LSTM networks using the slower update rates of $\Delta t = 3s$ and $\Delta t = 5s$, were 8.7 seconds and 10.7 seconds, respectively. The longer time longer time to accurately predict instability for the slower update rates of Δt can be attributed partly to a lower classification accuracy, and partly to the fact that measurements are being updated less frequently.

7.4.3 Generalization capability and training set requirement

The generalization capability of a ML method refers to the capability to generalize the learning from the actual training set to other, yet unseen, cases. Such capability is especially valuable in overcoming the combinatorial increase of complexity in the training when $N-1-1$ cases are also considered [75].

In Fig. 7.8, the classification accuracy is presented on the $N-1-1$ test set when the LSTM network have been trained on three different training sets. The results are presented when the network was trained on i) the full training set with all $N-1$




Figure 7.8: Classification accuracy over time when varying the number of $N-1-1$ cases included in the training data.

and $N-1-1$ cases included, ii) a smaller training set with all $N-1$ cases but where only a small batch (5,000) of $N-1-1$ cases have been included, and iii) a training set where the network is *only* trained on $N-1$. The same training approach as previously described were used. According to Fig. 7.8, the classification accuracy was significantly reduced when no $N-1-1$ cases are included in the test set. When including the small batch (5,000) of $N-1-1$ cases, the classification accuracy increased significantly. However, the accuracy is still lower than when the full training set is used. Thus, the importance of obtaining a representative training set is still imperative if a high classification accuracy is to be achieved.

7.5 Practical applications and requirements

The method is proposed to be used as an online tool for system operators to predict the system's future stability given the current state. It should be stressed that the method is not proposed to replace conventional voltage instability detection methods, but rather function as a supplementary tool to provide early warnings. The instantaneous prediction capability of the proposed method has to be weighed against the possibility of misclassification of the system's future stability. When comparing the proposed method to other conventional indicators for voltage instability detection (see [10]), it is important to remember that these might be more accurate once instability detected, but generally take *significantly* longer time to indicate instability, thus reducing the time frame that system operators have to steer the system back into stable operation.

The proposed method is mainly intended for predicting mid-term or long-term voltage instability where system operators will have the possibility to act on the warnings provided by the network. Theoretically, the method could be adapted to also handle short-term voltage instability. However, this would require more frequent measurement updates to ensure that the onset of short-term instability is detected in time. Because of the difference in the dynamical trajectories of the system for the two different types of instability events, training a separate LSTM network would likely provide better performance. Furthermore, the signals provided by the network would have to automatically trigger emergency controls, since the available time for system operators to act on the signals would be too short for manual control actions.

For the proposed method to be effective in prediction of long-term voltage instability, measurement updates should be available within a few seconds. Here, a measurement update rate each second have been assumed to be available. As was found in Section 7.4.2, slower measurement update rates lead both to lower classification accuracy and slower predictions. To assure that errors and missing values are filtered out, measurements should always be preceded by a state estimator. However, state estimates from a non-linear state estimator based on remote terminal units may be too slow to be effective. Thus, time-synchronized measurements from wide-area phasor measurements filtered through a linear state estimator would be preferred.

The softmax classifier of the LSTM network outputs a probability vector, where each class is given a certain probability. It should be noted that this probability vector does not provide a *true* representation of the model confidence. However, it can still be useful as a proxy by system operators to track the network's confidence in each prediction. Thus, the operator can use the probability vector directly in an online interface to track the network's belief in each prediction. Alternatively, argument maxima or other functions could be used to present the most probable prediction of the network, or, for instance, to avoid predictions of falsely labeled stable states.

The *practical* classification accuracy of the proposed method will be affected by many aspects and will generally be lower than on a simulated test set. One of the more important aspects are modeling errors, including erroneous system parameters or inaccurate modeling of parameter values for dynamic models. Such aspects will introduce a difference between the simulated and the actual dynamic response after a contingency. However, it should be noted that such limitations are not limited only to ML based approaches for VIP. All methods for DSA require that the dynamic models used in assessing the system response are accurately modeled.

CHAPTER 8

Conclusions and future research

This chapter summarizes the conclusions and refers back to the developed research questions and the aim of the thesis. Finally, a discussion of ideas and concepts for future research work is presented.

8.1 Conclusions

The main aim of this thesis has been to develop a new real-time voltage stability assessment tool (RVSAT) that can support system operators and allow more efficient utilization of the transmission grid. The thesis has demonstrated and tested two new methods for use in real-time voltage stability assessment.

In Section 1.2, four different research questions were identified based on the problem overview. The aim of the first research question was to examine the impact of higher penetration of power electronic controlled equipment on current methods for security margin assessment. This has been examined in the thesis through both theoretical and numerical analyses of the difference between the static and the dynamic voltage security margin. The difference between the DVSM and the conventional static VSM has been illustrated using a method called transient P - V curves, where the advantages of using the DVSM were established. The numerical comparison of the two methods confirms the advantages of the DVSM in power systems with high penetration of loads with fast load restoration dynamics. The results further indicated that the difference between the methods is smaller for cases when a fast fault-clearing time is possible.

In the second research question, the need for improvements in current methods for security margin estimation was emphasized. In response to this, a methodology for fast estimation of the DVSM has been developed with the aim to overcome the computational difficulties of computing the DVSM. The method uses a regression-based NN to provide a qualified guess of the actual DVSM. Moreover, a second NN is used to provide a classification of which contingency will be dimensioning for

the system. The estimates from the NNs are used in a method called dual binary search which is used to *validate* the actual DVSM using time-domain simulations. The ML-based approach is thus only proposed to *support* the estimation of the DVSM, while the actual DVSM is always validated through actual time-domain simulations. This two-step approach is proposed with the aim to overcome robustness issues and uncertainty of using ML-based methods, while still allowing near real-time estimations of the DVSM. The results are promising and the trained NNs provided good estimations of both the DVSM and classifications for the dimensioning contingency. The accurate estimations used in combination with the proposed dual binary search method was found to successfully reduce the required number of time-domain simulations, which would allow system operators to overcome the main practical difficulties of estimating the DVSM. These results also refer back to the fourth research question which aimed to identify practical barriers for system operators' implementations of ML-based methods.

The third research question aimed to develop methods that are both advanced enough to capture the intricate dynamics during a voltage collapse, while at the same time are fast enough to be used in real-time. This has been achieved in the thesis by developing a new approach for online VIP using an LSTM network capable of utilizing a sequence of measurements to improve classification accuracy. Once trained, the LSTM network can allow system operators to continuously assess and predict whether the present system state is stable or if it will evolve into an alert or an emergency state in the near future. The network is also adapted to be able to indicate *where* instability emerges, allowing system operators to perform more cost-effective control measures. The LSTM network was proposed to improve the available state signal by implicitly learning the dynamical trajectories of a power system following a disturbance. The LSTM architecture and the operation of the gating units ensure that the network is capable of capturing the long-term dependencies that are common in voltage instability events.

The results are encouraging and the proposed method is shown to have high accuracy in predicting voltage instability. Almost all $N-1$ contingency test cases were predicted correctly, and $N-1-1$ contingency test cases were predicted with over 95 % accuracy only seconds after a disturbance. The impact of the sequence length of the LSTM network was tested and indicated that a longer sequence provided a significantly better classification capability than both a feedforward NN and a network using a shorter sequence. Moreover, the generalization capability of the proposed LSTM network has been examined, where the classification accuracy on $N-1-1$ cases was assessed when the system was only trained on $N-1$ cases. It was found that this reduced the classification accuracy significantly, whereas including a smaller subset of $N-1-1$ cases into the training set resulted in significantly better performance.

8.2 Future research

Future research topics related to the conclusions of this thesis include the following:

- The developed methods for emergency and preventive monitoring in the RVSAT provide system operators with real-time warnings and estimations of security margins. Future research should include methods that can optimally act on this information. For instance, such methods could include an optimization-based protection scheme against voltage collapses. It would also be valuable to develop methods that can accurately determine what preventive control actions are most effective to increase the security margins in a power system for a certain operating point.
- The developed methods should also be tested on a real power system to examine the capacity and limitations. Regarding the proposed method for estimating the DVSM, it would be highly relevant to examine if the proposed method is capable of increasing the computational efficiency sufficiently to allow system operators to use it in real-time monitoring and operation. The developed methods for VIP is more difficult to test in a real system since larger disturbances cannot be applied without a large impact on the actual system. This method would instead have to be more rigorously tested through simulations. Finally, actual numerical comparisons between the DVSM and the VSM in real power systems would be of high interest, especially when considering high system penetration of loads with fast dynamic responses.
- More research should be assigned the uncertainties related to the developed ML algorithms. A difficulty in any supervised learning application is generating a representative training set. This is especially difficult for a large power system with a wide range of different operating conditions. Thus, to further examine the capability to generalize the learning from the actual training set to other, yet unseen, cases are of special interest. Furthermore, the sensitivity of the methods to various model and measurement errors should be examined in further detail.

Bibliography

- [1] IRENA, “Renewable power generation costs in 2018,” International Renewable Energy Agency, Abu Dabhi, Tech. Rep., 2019.
- [2] GWEC, “Global wind report 2018,” Global Wind Energy Council, Tech. Rep., Apr 2018, Available at: <https://gwec.net/global-wind-report-2018/>.
- [3] Energimyndigheten, “Energimyndigheten: Statistikdatabas,” Available at:<http://pxexternal.energimyndigheten.se/pxweb/sv/Vindkraftsstatistik/>, accessed: 2020-02-05.
- [4] S. Vindenergi, “Statistics and forecast,” Available at:<https://svenskvindenergi.org/statistik/2017-3>, accessed: 2020-03-18.
- [5] “Systemutvecklingsplan 2020–2029,” Statnett, Fingrid, Energinet, Svenska Kraftnät, Tech. Rep., Dec 2019, Available at: www.svk.se/om-oss/rapporter-och-remissvar.
- [6] ENTSO-E, “Picasso,” http://www.entsoe.eu/network_codes/eb/picasso/, accessed: 2020-02-06.
- [7] “Nordic grid development plan 2019,” Svenska Kraftnät, Tech. Rep., Jun 2019, Available at: www.svk.se/om-oss/rapporter-och-remissvar.
- [8] P. Kundur, J. Paserba, V. Ajjarapu, G. Andersson, A. Bose, C. Canizares, N. Hatziargyriou, D. Hill, A. Stankovic, C. Taylor, T. V. Cutsem, and V. Vittal, “Definition and classification of power system stability IEEE/CIGRE joint task force on stability terms and definitions,” *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1387–1401, Aug 2004.
- [9] T. Van Cutsem and C. Vournas, *Voltage stability of electric power systems*. Boston: Kluwer Academic Publishers, 1998.
- [10] M. Glavic and T. Van Cutsem, “A short survey of methods for voltage instability detection,” in *Proc. (IEEE) PES General Meeting*, Detroit, MI, Jul 2011, pp. 1–8.
- [11] T. Van Cutsem, C. Moisse, and R. Mailhot, “Determination of secure operating limits with respect to voltage collapse,” *IEEE Transactions on Power Systems*,

- vol. 14, no. 1, pp. 327–335, Feb 1999.
- [12] V. Ajjarapu and C. Christy, “The continuation power flow: a tool for steady state voltage stability analysis,” *IEEE Transactions on Power Systems*, vol. 7, no. 1, pp. 416–423, Feb 1992.
- [13] H.-D. Chiang *et al.*, “CPFLOW: a practical tool for tracing power system steady-state stationary behavior due to load and generation variations,” *IEEE Transactions on Power Systems*, vol. 10, no. 2, pp. 623–634, May 1995.
- [14] S. C. Savulescu, *Real-Time Stability in Power Systems. [electronic resource] : Techniques for Early Detection of the Risk of Blackout.*, ser. Power Electronics and Power Systems. Springer International Publishing, 2014.
- [15] I. Konstantelos, G. Jamgotchian, S. H. Tindemans, P. Duchesne, S. Cole, C. Merckx, G. Strbac, and P. Panciatici, “Implementation of a massively parallel dynamic security assessment platform for large-scale grids,” *IEEE Transactions on Smart Grid*, vol. 8, no. 3, pp. 1417–1426, May 2017.
- [16] T. Van Cutsem and R. Mailhot, “Validation of a fast voltage stability analysis method on the Hydro-Quebec system,” *IEEE Transactions on Power Systems*, vol. 12, no. 1, pp. 282–292, Feb 1997.
- [17] S. Brahma, R. Kavasseri, H. Cao, N. R. Chaudhuri, T. Alexopoulos, and Y. Cui, “Real-time identification of dynamic events in power systems using pmu data, and potential applications—models, promises, and challenges,” *IEEE Transactions on Power Delivery*, vol. 32, no. 1, pp. 294–301, 2017.
- [18] M. Glavic and T. Van Cutsem, “Wide-area detection of voltage instability from synchronized phasor measurements. part I: Principle,” *IEEE Transactions on Power Systems*, vol. 24, no. 3, pp. 1408–1416, 2009.
- [19] T. Van Cutsem, L. Wehenkel, M. Pavella, B. Heilbronn, and M. Goubin, “Decision tree approaches to voltage security assessment,” *IEE Proceedings C - Generation, Transmission and Distribution*, vol. 140, no. 3, pp. 189–198, May 1993.
- [20] Y. Mansour, A. Y. Chang, J. Tamby, E. Vaahedi, B. R. Corns, and M. A. El-Sharkawi, “Large scale dynamic security screening and ranking using neural networks,” *IEEE Transactions on Power Systems*, vol. 12, no. 2, pp. 954–960, May 1997.
- [21] K. Sun, S. Likhate, V. Vittal, V. S. Kolluri, and S. Mandal, “An online dynamic security assessment scheme using phasor measurements and decision trees,” *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1935–1943, Nov 2007.
- [22] H. Khoshkhou and S. M. Shahrtash, “Fast online dynamic voltage instability prediction and voltage stability classification,” *IET Generation, Transmission & Distribution*, vol. 8, no. 5, pp. 957–965, May 2014.

-
- [23] C. Liu, F. Tang, and C. L. Bak, “An accurate online dynamic security assessment scheme based on random forest,” *Energies*, vol. 11, no. 7, 2018.
- [24] T. Van Cutsem, M.-E. Grenier, and D. Lefebvre, “Combined detailed and quasi steady-state time simulations for large-disturbance analysis,” *International Journal of Electrical Power & Energy Systems*, vol. 28, no. 9, pp. 634 – 642, 2006.
- [25] A. Sittithumwat and K. Tomsovic, “Dynamic security margin estimation using artificial neural networks,” in *IEEE Power Engineering Society Summer Meeting*, vol. 3, Chicago, IL, July 2002, pp. 1322–1327.
- [26] N. Amjady, “Dynamic voltage security assessment by a neural network based method,” *Electric Power Systems Research*, vol. 66, no. 3, pp. 215 – 226, 2003.
- [27] J. Shang, J. Zhang, W. Zhao, and J. Liu, “ANN based dynamic voltage security assessment for a practical power system,” in *2007 International Power Engineering Conference (IPEC 2007)*, Singapore, Dec 2007, pp. 794–798.
- [28] M. V. Baghmisheh and H. Razmi, “Dynamic voltage stability assessment of power transmission systems using neural networks,” *Energy Conversion and Management*, vol. 49, no. 1, pp. 1 – 7, 2007.
- [29] R. F. Nuqui, “State estimation and voltage security monitoring using synchronized phasor measurements,” Ph.D. dissertation, Virginia Tech, 2001.
- [30] R. Diao, K. Sun, V. Vittal, R. J. O’Keefe, M. R. Richardson, N. Bhatt, D. Stradford, and S. K. Sarawgi, “Decision tree-based online voltage security assessment using PMU measurements,” *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 832–839, May 2009.
- [31] H. Khoshkhoo and S. M. Shahrtash, “On-line dynamic voltage instability prediction based on decision tree supported by a wide-area measurement system,” *IET Generation, Transmission & Distribution*, vol. 6, no. 11, pp. 1143–1152, November 2012.
- [32] M. L. Scala, M. Trovato, and F. Torelli, “A neural network-based method for voltage security monitoring,” *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1332–1341, Aug 1996.
- [33] P. Geurts and L. Wehenkel, “Early prediction of electric power system blackouts by temporal machine learning,” in *Proceedings of the ICML98/AAAI98 Workshop on “Predicting the future: AI Approaches to time series analysis*, 1998, pp. 24–26.
- [34] L. Wehenkel, C. Lebrevelec, M. Trotignon, and J. Batut, “Probabilistic design of power-system special stability controls,” *Control Engineering Practice*, vol. 7, no. 2, pp. 183–194, 1999.
- [35] P. Geurts and L. Wehenkel, “Temporal machine learning for switching con-

- trol,” in *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer, 2000, pp. 401–408.
- [36] P. Kundur, J. Paserba, V. Ajjarapu, G. Andersson, A. Bose, C. Canizares, N. Hatziargyriou, D. Hill, A. Stankovic, C. Taylor, T. V. Cutsem, and V. Vittal, “Definition and classification of power system stability ieee/cigre joint task force on stability terms and definitions,” *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1387–1401, Aug 2004.
- [37] K. Morison, Lei Wang, and P. Kundur, “Power system security assessment,” *IEEE Power and Energy Magazine*, vol. 2, no. 5, pp. 30–39, Sep. 2004.
- [38] ETSO: European Transmission System Operators, “Definitions of transfer capacities in liberalised electricity markets,” https://www.entsoe.eu/fileadmin/user_upload/_library/ntc/entsoe_transferCapacityDefinitions.pdf, accessed: 2020-02-18.
- [39] S. Yari and H. Khoshkhoo, “A comprehensive assessment to propose an improved line stability index,” *International Transactions on Electrical Energy Systems*, vol. 29, no. 4, p. e2806, 2019.
- [40] J. Modarresi, E. Gholipour, and A. Khodabakhshian, “A comprehensive review of the voltage stability indices,” *Renewable and Sustainable Energy Reviews*, vol. 63, pp. 1–12, 2016.
- [41] K. Vu, M. M. Begovic, D. Novosel, and M. M. Saha, “Use of local measurements to estimate voltage-stability margin,” *IEEE Transactions on Power Systems*, vol. 14, no. 3, pp. 1029–1035, 1999.
- [42] A. Perez, H. Jóhannsson, J. Østergaard, M. Glavic, and T. Van Cutsem, “Improved thévenin equivalent methods for real-time voltage stability assessment,” in *2016 IEEE International Energy Conference (ENERGYCON)*, 2016, pp. 1–6.
- [43] Y. Wang, I. R. Pordanjani, W. Li, W. Xu, T. Chen, E. Vaahedi, and J. Gurney, “Voltage stability monitoring based on the concept of coupled single-port circuit,” *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2154–2163, 2011.
- [44] F. Gubina and B. Strmcnik, “Voltage collapse proximity index determination using voltage phasors approach,” *IEEE Transactions on Power Systems*, vol. 10, no. 2, pp. 788–794, 1995.
- [45] M. Moghavvemi and O. Faruque, “Real-time contingency evaluation and ranking technique,” *IEEE Proc. - Generation, Transmission and Distribution*, vol. 145, no. 5, pp. 517–524, 1998.
- [46] M. Moghavvemi and F. M. Omar, “Technique for contingency monitoring and voltage collapse prediction,” *IEE Proc. - Generation, Transmission and Distribution*, vol. 145, no. 6, pp. 634–640, 1998.

-
- [47] C. Vournas and T. V. Cutsem, “Local identification of voltage emergency situations,” *IEEE Transactions on Power Systems*, vol. 23, pp. 1239 – 1248, 09 2008.
- [48] C. D. Vournas, C. Lambrou, and M. Kanatas, “Application of local autonomous protection against voltage instability to iee test system,” *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 3300–3308, July 2016.
- [49] C. Vournas, C. Lambrou, and P. Mandoulidis, “Voltage stability monitoring from a transmission bus pmu,” in *2017 IEEE Manchester PowerTech*, June 2017, pp. 1–1.
- [50] C. Vournas and P. Mandoulidis, “On-line voltage stability monitoring,” in *Circuits and Systems (ISCAS), 2018 IEEE Int. Symposium on*. IEEE, 2018, pp. 1–5.
- [51] J. Liu and C. Chu, “Wide-area measurement-based voltage stability indicators by modified coupled single-port models,” *IEEE Transactions on Power Systems*, vol. 29, no. 2, pp. 756–764, March 2014.
- [52] B. Cui and Z. Wang, “Voltage stability assessment based on improved coupled single-port method,” *IET Generation, Transmission Distribution*, vol. 11, no. 10, pp. 2703–2711, 2017.
- [53] Lixin Bao, Zhenyu Huang, and Wilsun Xu, “Online voltage stability monitoring using var reserves,” *IEEE Transactions on Power Systems*, vol. 18, no. 4, pp. 1461–1469, Nov 2003.
- [54] B. Leonardi and V. Ajjarapu, “Investigation of various generator reactive power reserve (grpr) definitions for online voltage stability/security assessment,” in *2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, July 2008, pp. 1–7.
- [55] —, “An approach for real time voltage stability margin control via reactive power reserve sensitivities,” *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 615–625, May 2013.
- [56] H. Nguyen Duc, I. Kamwa, L.-A. Dessaint, and H. Cao-Duc, “A novel approach for early detection of impending voltage collapse events based on the support vector machine,” *International Transactions on Electrical Energy Systems*, vol. 27, no. 9, p. e2375, 2017.
- [57] M. Glavic and T. Van Cutsem, “Wide-area detection of voltage instability from synchronized phasor measurements. part II: Simulation results,” *IEEE Transactions on Power Systems*, vol. 24, no. 3, pp. 1417–1425, 2009.
- [58] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [59] K. Hornik, “Approximation capabilities of multilayer feedforward networks,”

- Neural Networks*, vol. 4, no. 2, pp. 251 – 257, 1991. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/089360809190009T>
- [60] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, no. 6088, p. 533, 1986.
- [61] P. J. Werbos *et al.*, “Backpropagation through time: what it does and how to do it,” *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [62] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber *et al.*, “Gradient flow in recurrent nets: the difficulty of learning long-term dependencies,” 2001.
- [63] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [64] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, “LSTM: A Search Space Odyssey,” *IEEE Transactions on Neural Networks & Learning Systems*, vol. 28, no. 10, pp. 2222–2232, Oct 2017.
- [65] T. Van Cutsem, M. Glavic, W. Rosehart, and J. A. dos Santos, “Test systems for voltage stability analysis and security assessment,” IEEE/PES Task Force, Tech. Rep. PES-TR19, Aug. 2015. [Online]. Available: <http://resourcecenter.ieee-pes.org/pes/product/technical-publications/PESTR19>
- [66] M. K. Pal, “Voltage stability conditions considering load characteristics,” *IEEE Transactions on Power Systems*, vol. 7, no. 1, pp. 243–249, Feb 1992.
- [67] I. Dobson, “The irrelevance of load dynamics for the loading margin to voltage collapse and its sensitivities,” in *Bulk power system voltage phenomena - III: Voltage stability, security & control*, 1994, pp. 509–518.
- [68] C. Dharmakeerthi, N. Mithulananthan, and T. Saha, “Impact of electric vehicle fast charging on power system voltage stability,” *International Journal of Electrical Power & Energy Systems*, vol. 57, pp. 241 – 249, 2014.
- [69] *PSS®E 34.2.0 Model Library*, Siemens Power Technologies International, Schenectady, NY, Apr. 2017.
- [70] D. Q. Zhou, U. D. Annakkage, and A. D. Rajapakse, “Online monitoring of voltage stability margin using an artificial neural network,” *IEEE Transactions on Power Systems*, vol. 25, no. 3, pp. 1566–1574, Aug 2010.
- [71] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv e-prints*, p. arXiv:1412.6980, Dec 2014.
- [72] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, Massachusetts London, England: The MIT Press, 2015.
- [73] H. Hagmar, L. A. Tuan, O. Carlson, and R. Eriksson, “On-line voltage instability prediction using an artificial neural network,” in *2019 IEEE Milano PowerTech*, June 2019 (Accepted), pp. 1–6.

- [74] F. Thams, A. Venzke, R. Eriksson, and S. Chatzivasileiadis, “Efficient database generation for data-driven security assessment of power systems,” *IEEE Transactions on Power Systems*, pp. 1–1, 2019.
- [75] P. Mitra, V. Vittal, B. Keel, and J. Mistry, “A systematic approach to $n-1-1$ analysis for power system security assessment,” *IEEE Power and Energy Technology Systems Journal*, vol. 3, no. 2, pp. 71–80, June 2016.

APPENDIX A

Dynamic parameters for the IEESGO model

This appendix presents the model and the dynamic parameters used for the IEESGO model. The IEESGO takes an active power reference as inputs which allows controlling the mechanical power of the machine directly. Only a very small governor lag have been used for this specific model.

Table A.1: Results for Case A: Three-phased fault followed by tripping a line.

Parameter	Value	Description
T_1	0.0	Controller lag (sec)
T_2	0.0	Controller lead compensation (sec)
T_3	0.05	Governor lag (sec)
T_4	0.0	Delay due to steam inlet (sec)
T_5	0.0	Reheater delay (sec)
T_6	0.0	Delay due to IP-LP turbine, crossover pipes, and LP end hoods (sec)
K_1	0.0	1/per unit regulation
K_2	0.0	Fraction
K_3	0.0	Fraction
P_{MAX}	1.0	Upper power limit
P_{MIN}	0.0	Lower power limit

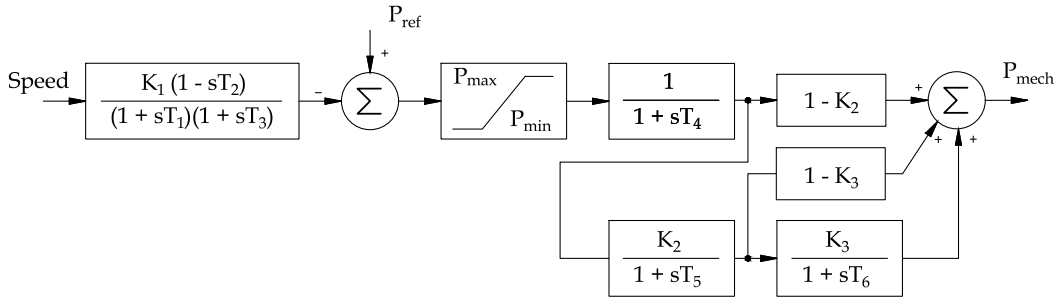


Figure A.1: Block diagram of the IEESGO Model.