Neural Networks For Time Series Forecasting Practical

Practical Time Series Analysis
Forecasting: principles and practice
Artificial Neural Networks and Time Series Models in Economic Forecasting
Predict the Future with MLPs, CNNs and LSTMs in Python
Time Series Analysis Using Neural Networks
Recurrent Neural Networks for Time Series Prediction
Time Series Forecasting using Deep Learning
Time Series Forecasting Using Neural Networks
State-of-the-art Deep Learning for Multi-product Intermittent Time Series Forecasting
Informatics and Cybernetics in Intelligent Systems
23rd International Conference on Artificial Neural Networks, Sofia, Bulgaria, September 10-13, 2013, Proceedings
Case Studies in Economics
from Training to Prediction
Nonlinear Time Series Forecasting with Neural Networks
Using Artificial Neural Networks for Timeseries Smoothing and Forecasting
Neural Networks in Business Forecasting
Neural Networks in Time Series Forecasting
Over 50 recipes for applying modern Python libraries to financial data analysis
Theoretical and Applied Mathematics in International Business
Artificial neural networks in time series forecasting
Neural Networks for Time Series Analysis
An Intuitive Introduction to Deep Learning for Applied Time Series Modeling
Forecast, predict, and detect anomalies with state-of-the-art machine learning methods
Using Neural Networks to Do Time Series Forecasting
An Overview and Comparative Analysis
Neural Networks for Time Series Forecasting with R
10th International Conference, ICIC 2014, Taiyuan, China, August 3-6, 2014, Proceedings
Intelligent Computing in Bioinformatics
Neural Networks for Time Series Forecasting
Forecasting of Financial Markets
Deep Learning for Time Series Forecasting
Principles of Forecasting
Long Short-Term Memory Networks With Python
Stochastic Neural Networks and Their Applications to Regression Analysis and Time Series Forecasting
Neural Networks for Seasonal Time Series Forecasting
Proceedings of 10th Computer Science On-line Conference 2021, Vol. 3
Develop Sequence Prediction Models with Deep Learning
An Intuitive Step by Step Blueprint for Beginners

Combining PyTorch, RNN, TCN, and Deep Neural Network Models to Provide Production-Ready Prediction Solutions (English Edition)

ALBERT JAYLEEN

Practical Time Series Analysis Springer
The analysis of a time series is a problem well known to statisticians. Neural networks form the basis of an entirely non-linear approach to the analysis of time series. It has been widely used in pattern recognition, classification and prediction. Recently, reviews from a statistical perspective were done by Cheng and Titterington (1994) and Ripley (1993). One of the most important properties of a neural network is its ability to learn. In neural network methodology, the data set is divided in three different sets, namely a training set, a cross-validation set, and a test set. The training set is used for training the network with the various available learning (optimisation) algorithms. Different algorithms will perform best on different problems. The advantages and limitations of different algorithms in respect of all training problems are discussed. In this dissertation the method of neural networks and that of ARIMA. models are discussed. The procedures of identification, estimation and evaluation of both models are investigated. Many of the standard techniques in statistics can be compared with neural network methodology, especially in applications with large data sets. Additional information available on two discs stored at the Africana section, Merensky Library.
Forecasting: principles and practice Packt Publishing Ltd
Deep learning methods offer a lot of promise for time series forecasting, such as the automatic learning of temporal dependence and the automatic handling of temporal structures like trends and seasonality. With clear explanations, standard Python libraries, and step-by-step tutorial lessons you'll discover how to develop deep learning models for your own time series forecasting projects.
Artificial Neural Networks and Time Series Models in Economic Forecasting BoD – Books on Demand
This book – in conjunction with the volumes LNCS 8588 and LNAI 8589 – constitutes the refereed proceedings of the 10th International Conference on Intelligent Computing, ICIC 2014, held in Taiyuan, China, in August 2014. The 58 papers of this volume were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections such as machine learning; neural networks; image processing; computational systems biology and medical informatics; biomedical informatics theory and methods; advances on bio-inspired computing; protein and gene bioinformatics: analysis, algorithms, applications.
Predict the Future with MLPs, CNNs and LSTMs in Python Springer Nature
Explore the infinite possibilities offered by Artificial Intelligence and Neural Networks KEY FEATURES ● Covers numerous concepts, techniques, best practices and troubleshooting tips by community experts. ● Includes practical demonstration of robust deep learning prediction models with exciting use-cases. ● Covers the use of the most powerful research toolkit such as Python, PyTorch, and Neural Network Intelligence. DESCRIPTION This book is amid at teaching the readers how to apply the deep learning techniques to the time series forecasting challenges and how to build prediction models using PyTorch. The readers will learn the fundamentals of PyTorch in the early stages of the book. Next, the time series forecasting is covered in greater depth after the programme has been developed. You will try to use machine learning to identify the patterns that can help us forecast the future results. It covers methodologies such as Recurrent Neural Network, Encoder-decoder model, and Temporal Convolutional Network, all of which are state-of-the-art neural network architectures. Furthermore, for good measure, we have also introduced the neural architecture search, which automates searching for an ideal neural network design for a certain task. Finally by the end of the book, readers would be able to solve complex real-world prediction issues by applying the models and strategies learnt throughout the course of the book. This book also offers another great way of mastering deep learning and its various techniques. WHAT YOU WILL LEARN ● Work with the Encoder-Decoder concept and Temporal Convolutional Network mechanics. ● Learn the basics of neural architecture search with Neural Network Intelligence. ● Combine standard statistical analysis methods with deep learning approaches. ● Automate the search for optimal predictive architecture. ● Design your custom neural network architecture for specific tasks. ● Apply predictive models to real-world problems of forecasting stock quotes, weather, and natural processes. WHO THIS BOOK IS FOR This book is written for engineers, data scientists, and stock traders who want to build time series forecasting programs using deep learning. Possessing some familiarity of Python is sufficient, while a basic understanding of machine learning is desirable but not needed. TABLE OF CONTENTS 1. Time Series Problems and Challenges 2. Deep Learning with PyTorch 3. Time Series as Deep Learning Problem 4. Recurrent Neural Networks 5. Advanced Forecasting Models 6. PyTorch Model Tuning with Neural Network Intelligence 7. Applying Deep Learning to Real-world Forecasting Problems 8. PyTorch Forecasting Package 9. What is Next?
Time Series Analysis Using Neural Networks Springer
The book constitutes the proceedings of the 23rd International Conference on Artificial Neural Networks, ICANN 2013, held in Sofia, Bulgaria, in September 2013. The 78 papers included in the proceedings were carefully reviewed and selected from 128 submissions. The focus of the papers is on following topics: neurofinance graphical network models, brain machine interfaces, evolutionary neural networks, neurodynamics, complex systems, neuroinformatics, neuroengineering, hybrid systems, computational biology, neural hardware, bioinspired embedded systems, and collective intelligence.
Recurrent Neural Networks for Time Series Prediction IGI Global
The Long Short-Term Memory network, or LSTM for short, is a type of recurrent neural network that achieves state-of-the-art results on challenging prediction problems. In this laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about LSTMs. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what LSTMs are, and how to develop a suite of LSTM models to get the most out of the method on your sequence prediction problems.
Time Series Forecasting using Deep Learning Time Series Forecasting using Deep LearningCombining PyTorch, RNN, TCN, and Deep Neural Network Models to Provide Production-Ready Prediction Solutions (English Edition)
Esta dissertação investiga a utilização de Redes Neurais Artificiais (RNAs) na previsão de séries

temporais, em particular de séries financeiras, consideradas uma classe especial de séries temporais, caracteristicamente ruídos e sem periodicidade aparente. O trabalho envolve quatro partes principais: um estudo sobre redes neurais artificiais e séries temporais; a modelagem das RNAs para previsão de séries temporais; o desenvolvimento de um ambiente de simulação; e o estudo de caso. No estudo sobre Redes Neurais Artificiais e séries temporais fez-se um levantamento preliminar das aplicações de RNAs na previsão de séries. Constatou-se a predominância do uso do algoritmos de retropropagação do erro para o treinamento das redes, bem como dos modelos estatísticos de regressão, de médias móveis e de alisamento exponencial nas comparações com os resultados da rede. Na modelagem das RNAs de retropropagação do erro considerou-se três fatores determinantes no desempenho da rede: convergência, generalização e escalabilidade. Para o controle destes fatores usou-se mecanismos como; escolha da função de ativação dos neurônios sigmóide ou tangente hiperbólica; escolha da função erro MSE (Mean Square Error) ou MAD (Mean Absolutd Deviation); e escolha dos parâmetros de controle do gradiente descendente e do temapo de treinamento taxa de aprendizado e termo de momento. Por fim, definiu-se a arquitetura da rede em função da técnica utilizada para a identificação de regularidades na série (windowing) e da otimização dos fatores indicadores de desempenho da rede. O ambiente de simulação foi desenvolvido em linguagem C e contém 3.600 linhas de códigos divididas em três módulos principais: interface com o usuário, simulação e funções secundárias. O módulo de interface com o usuário é responsável pela configuração e parametrização da rede, como também pela visualização gráfica dos resultados; módulo de simulação executa as fases de treinamento e testes das RNAs; o módulo de funções secundárias cuida do pré/pós-processamento dos dados, da manipulação de arquivos e dos cálculos dos métodos de avaliação empregados. No estudo de caso, as RNAs foram modeladas para fazer previsões da série do preço do ouro no mercado internacional. Foram feitas previsões univariadas single e multi-step e previsões multivariadas utilizando taxas de câmbio de moedas estrangeiras. Os métodos utilizandos para a avaliação do desempenho da rede foram: coeficiente U de Theil, MSE (Mean Square Error), NRMSE (Normalized Root Mean Square Error), POCID (Percentage Of Change In Direction), scattergram e comparação gráfica. Os resultados obtidos, além de avaliados com os métodos acima, foram comparados com o modelo de Box-Jenkins e comprovaram a superioridade das RNAs no tratamento de dados não-lineares e altamente ruidosos.

Time Series Forecasting Using Neural Networks Simon and Schuster

The aim of this publication is to identify and apply suitable methods for analysing and predicting the time series of gold prices, together with acquainting the reader with the history and characteristics of the methods and with the time series issues in general. Both statistical and econometric methods, and especially artificial intelligence methods, are used in the case studies. The publication presents both traditional and innovative methods on the theoretical level, always accompanied by a case study, i.e. their specific use in practice. Furthermore, a comprehensive comparative analysis of the individual methods is provided. The book is intended for readers from the ranks of academic staff, students of universities of economics, but also the scientists and practitioners dealing with the time series prediction. From the point of view of practical application, it could provide useful information for speculators and traders on financial markets, especially the commodity markets.

State-of-the-art Deep Learning for Multi-product Intermittent Time Series Forecasting Springer Science & Business Media

The key component in forecasting demand and consumption of resources in a supply network is an accurate prediction of real-valued time series. Indeed, both service interruptions and resource waste can be reduced with the implementation of an effective forecasting system. Significant research has thus been devoted to the design and development of methodologies for short term load forecasting over the past decades. A class of mathematical models, called Recurrent Neural Networks, are nowadays gaining renewed interest among researchers and they are replacing many practical implementations of the forecasting systems, previously based on static methods. Despite the undeniable expressive power of these architectures, their recurrent nature complicates their understanding and poses challenges in the training procedures. Recently, new important families of recurrent architectures have emerged and their applicability in the context of load forecasting has not been investigated completely yet. This work performs a comparative study on the problem of Short-Term Load Forecast, by using different classes of state-of-the-art Recurrent Neural Networks. The authors test the reviewed models first on controlled synthetic tasks and then on different real datasets, covering important practical cases of study. The text also provides a general overview of the most important architectures and defines guidelines for configuring the recurrent networks to predict real-valued time series.

Informatics and Cybernetics in Intelligent Systems BPB Publications

Learn how to apply the principles of machine learning to time series modeling with this indispensable resource Machine Learning for Time Series Forecasting with Python is an incisive and straightforward examination of one of the most crucial elements of decision-making in finance, marketing, education, and healthcare: time series modeling. Despite the centrality of time series forecasting, few business analysts are familiar with the power or utility of applying machine learning to time series modeling. Author Francesca Lazzeri, a distinguished machine learning scientist and economist, corrects that deficiency by providing readers with comprehensive and approachable explanation and treatment of the application of machine learning to time series forecasting. Written for readers who have little to no experience in time series forecasting or machine learning, the book comprehensively covers all the topics necessary to: Understand time series forecasting concepts, such as stationarity, horizon, trend, and seasonality Prepare time series data for modeling Evaluate time series forecasting models' performance and accuracy Understand when to use neural networks instead of traditional time series models in time series forecasting Machine Learning for Time Series Forecasting with Python is full real-world examples, resources and concrete strategies to help readers explore and transform data and develop usable, practical time series forecasts. Perfect for entry-level data scientists, business analysts, developers, and researchers, this book is an invaluable and indispensable guide to the fundamental and advanced concepts of machine learning applied to time series modeling.

23rd International Conference on Artificial Neural Networks, Sofia, Bulgaria, September 10-13, 2013, Proceedings Springer

Forecasting is required in many situations. Stocking an inventory may require forecasts of demand months in advance. Telecommunication routing requires traffic forecasts a few minutes ahead.

Whatever the circumstances or time horizons involved, forecasting is an important aid in effective and efficient planning. This textbook provides a comprehensive introduction to forecasting methods and presents enough information about each method for readers to use them sensibly.

Case Studies in Economics John Wiley & Sons

Step by Step guide filled with real world practical examples. About This Book Get your first experience with data analysis with one of the most powerful types of analysis—time-series. Find patterns in your data and predict the future pattern based on historical data. Learn the statistics, theory, and implementation of Time-series methods using this example-rich guide Who This Book Is For This book is for anyone who wants to analyze data over time and/or frequency. A statistical background is necessary to quickly learn the analysis methods. What You Will Learn Understand the basic concepts of Time Series Analysis and appreciate its importance for the success of a data science project Develop an understanding of loading, exploring, and visualizing time-series data Explore auto-correlation and gain knowledge of statistical techniques to deal with non-stationarity time series Take advantage of exponential smoothing to tackle noise in time series data Learn how to use auto-regressive models to make predictions using time-series data Build predictive models on time series using techniques based on auto-regressive moving averages Discover recent advancements in deep learning to build accurate forecasting models for time series Gain familiarity with the basics of Python as a powerful yet simple to write programming language In Detail Time Series Analysis allows us to analyze data which is generated over a period of time and has sequential interdependencies between the observations. This book describes special mathematical tricks and techniques which are geared towards exploring the internal structures of time series data and generating powerful descriptive and predictive insights. Also, the book is full of real-life examples of time series and their analyses using cutting-edge solutions developed in Python. The book starts with descriptive analysis to create insightful visualizations of internal structures such as trend, seasonality and autocorrelation. Next, the statistical methods of dealing with autocorrelation and non-stationary time series are described. This is followed by exponential smoothing to produce meaningful insights from noisy time series data. At this point, we shift focus towards predictive analysis and introduce autoregressive models such as ARMA and ARIMA for time series forecasting. Later, powerful deep learning methods are presented, to develop accurate forecasting models for complex time series, and under the availability of little domain knowledge. All the topics are illustrated with real-life problem scenarios and their solutions by best-practice implementations in Python. The book concludes with the Appendix, with a brief discussion of programming and solving data science problems using Python. Style and approach This book takes the readers from the basic to advance level of Time series analysis in a very practical and real world use cases.

from Training to Prediction Machine Learning Mastery

Machine learning uses two types of techniques: supervised learning, which trains a model on known input and output data so that it can predict future outputs, and unsupervised learning, which finds hidden patterns or intrinsic structures in input data. The aim of supervised machine learning is to build a model that makes predictions based on evidence in the presence of uncertainty. A supervised learning algorithm takes a known set of input data and known responses to the data (output) and trains a model to generate reasonable predictions for the response to new data. Supervised learning uses classification and regression techniques to develop predictive models. • Classification techniques predict categorical responses, for example, whether an email is genuine or spam, or whether a tumor is cancerous or benign. Classification models classify input data into categories. Typical applications include medical imaging, image and speech recognition, and credit scoring. • Regression techniques predict continuous responses, for example, changes in temperature or fluctuations in power demand. Typical applications include electricity load forecasting and algorithmic trading. This book develops time series forecasting techniques using neural networks

Nonlinear Time Series Forecasting with Neural Networks John Wiley & Sons

In the past, practical applications motivated the development of mathematical theories, which then became the subject of study in pure mathematics where abstract concepts are studied for their own sake. The activity of applied mathematics is thus intimately connected with research in pure mathematics, which is also referred to as theoretical mathematics. Theoretical and Applied Mathematics in International Business is an essential research publication that explores the importance and implications of applied and theoretical mathematics within international business, including areas such as finance, general management, sales and marketing, and supply chain management. Highlighting topics such as data mining, global economics, and general management, this publication is ideal for scholars, specialists, managers, corporate professionals, researchers, and academicians.

Using Artificial Neural Networks for Timeseries Smoothing and Forecasting OTexts

Comprehensively specified benchmarks are provided (including weight values), drawn from time series examples in chaos theory and financial futures. The book covers data preprocessing, random walk theory, trading systems and risk analysis. It also provides a literature review, a tutorial on backpropagation, and a chapter on further reading and software.

Neural Networks in Business Forecasting Packt Publishing Ltd

Artificial neural networks are suitable for many tasks in pattern recognition and machine learning. Unlike conventional techniques for time series analysis, an artificial neural network needs little information about the time series data and can be applied to a broad range of problems. The usage of artificial neural networks for time series analysis relies purely on the data that were observed. As Radial Basis networks with one hidden layer is capable of approximating any measurable function. An artificial neural network is powerful enough to represent any form of time series. The capability to generalize allows artificial neural networks to learn even in the case of noisy and/or missing data. Another advantage over linear models is the network's ability to represent nonlinear time series. Prediction of tides is very much essential for human activities and to reduce the construction cost in marine environment. This book presents an application of the artificial neural network with Radial basis function for accurate prediction of tides. This neural network model predicts the time series data of hourly tides directly while using an an efficient learning process.

Neural Networks in Time Series Forecasting CRC Press

The use of neural networks is permeating every area of signal processing. They can provide powerful means for solving many problems, especially in nonlinear, real-time, adaptive, and blind signal processing. The Handbook of Neural Network Signal Processing brings together applications that

were previously scattered among various publications to provide an up-to-date, detailed treatment of the subject from an engineering point of view. The authors cover basic principles, modeling, algorithms, architectures, implementation procedures, and well-designed simulation examples of audio, video, speech, communication, geophysical, sonar, radar, medical, and many other signals. The subject of neural networks and their application to signal processing is constantly improving. You need a handy reference that will inform you of current applications in this new area. The Handbook of Neural Network Signal Processing provides this much needed service for all engineers and scientists in the field.

Over 50 recipes for applying modern Python libraries to financial data analysis LAP Lambert Academic Publishing

Artificial intelligence (AI) is everywhere and it's here to stay. Most aspects of our lives are now touched by artificial intelligence in one way or another, from deciding what books or flights to buy online to whether our job applications are successful, whether we receive a bank loan, and even what treatment we receive for cancer. Artificial Neural Networks (ANNs) as a part of AI maintains the capacity to solve problems such as regression and classification with high levels of accuracy. This book aims to discuss the usage of ANNs for optimal solving of time series applications and clustering. Bounding of optimization methods particularly metaheuristics considered as global optimizers with ANNs make a strong and reliable prediction tool for handling real-life application. This book also demonstrates how different fields of studies utilize ANNs proving its wide reach and relevance.

Theoretical and Applied Mathematics in International Business Packt Publishing Ltd

Deep learning is gaining traction and considerable attention due to the state-of-the-art results obtained in computer vision, object detection, natural language processing, sequential analysis, and multiple other domains. Study of literature reveals that time series analysis is a good candidate for modeling using deep learning techniques. Time series analysis has applications from finance to supply chain domains and proves to be critical in driving organizations' profit and strategic growth. In a retail setting, product demand forecasting helps in minimizing inventory, optimizing service levels, and maximizing revenue. When dealing with demand forecasting, a much complex branch of intermittent demand profiles arises. When forecasting time series, the standard option comes down to statistical learning methods such as ARIMA, exponential smoothing, and several other models. However, in case of intermittency in demand and forecasting multiple time series at once, statistical learning methods fail to provide a high level of accuracy and can sometimes become computationally expensive as well. Deep learning algorithms enter the fray, as they can be applied to tackle the problem of forecasting intermittent sales while solving the problem in a computationally frugal manner. The study focuses on solving these two problems using a state-of-the-art based approach. It helps us answer the questions of -- How to implement neural networks in a value-add manner? And which models and architectures work best in our time series prediction problem with similar real-world applications? The study reveals that recurrent and convolutional architectures exhibit versatility and value in solving this problem, helping us understand the deep learning models and their application architectures in real-world scenarios. In this thesis, we have tried to answer these two important questions. The data was obtained from Kaggle for the M5 forecasting competition. The dataset relates to the daily Walmart sales of 3,000 products ranging across 10 stores. The data comprises of 3 different categories and 7 sub-categories, making it a multi-time series forecasting problem. We have applied the methods of statistical learning and deep learning to solve this problem. Statistical models of naïve method, moving average, ARIMA, Croston forecasting have been implemented. In deep learning, we initially use the deep feed-forward neural network to forecast the sales. Then recurrent architectures of RNN, LSTM and GRU are applied. Sequence learning and Attention mechanism have been implemented. Convolutional architectures of CNN, Wavenet, and temporal convolutional network have also been experimented for our problem. For the methodology, we initially select a single time series from the dataset and apply the statistical and deep learning models. This step in the methodology provides us with a strong fundamental understanding of how deep learning models are tuned to obtain the optimal architecture. Then, using the results from a single time series forecasting problem, we shortlist the most optimal deep learning models and their optimal architectures, to solve the problem of time series forecasting. We conclude that recurrent architectures provide the optimal solutions for our analysis (we define optimality through error minimization), and state-of-the-art models such as attention mechanism and sequence learning provide results within acceptable range, but their models are too computationally expensive to learn for multiple epochs and forecasts. We then conclude our analysis by providing important areas to focus on deep learning for time series forecasting in our future work.

Artificial neural networks in time series forecasting CRC Press

Time Series Forecasting using Deep LearningCombining PyTorch, RNN, TCN, and Deep Neural Network Models to Provide Production-Ready Prediction Solutions (English Edition)BPB Publications

Related with Neural Networks For Time Series Forecasting Practical:

• Comparison Bars In Math : click here