

1. Record Nr.	UNINA9910488729003321
Autore	Korstanje Joos
Titolo	Advanced Forecasting with Python : With State-Of-the-Art-Models Including LSTMs, Facebook's Prophet, and Amazon's DeepAR
Pubbl/distr/stampa	Berkeley, CA : , : Apress L. P., , 2021 ©2021
ISBN	1-4842-7150-5
Descrizione fisica	1 online resource (294 pages)
Disciplina	006.31
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Table of Contents -- About the Author -- About the Technical Reviewer -- Introduction -- Part I: Machine Learning for Forecasting -- Chapter 1: Models for Forecasting -- Reading Guide for This Book -- Machine Learning Landscape -- Univariate Time Series Models -- A Quick Example of the Time Series Approach -- Supervised Machine Learning Models -- A Quick Example of the Supervised Machine Learning Approach -- Correlation Coefficient -- Other Distinctions in Machine Learning Models -- Supervised vs. Unsupervised Models -- Classification vs. Regression Models -- Univariate vs. Multivariate Models -- Key Takeaways -- Chapter 2: Model Evaluation for Forecasting -- Evaluation with an Example Forecast -- Model Quality Metrics -- Metric 1: MSE -- Metric 2: RMSE -- Metric 3: MAE -- Metric 4: MAPE -- Metric 5: R2 -- Model Evaluation Strategies -- Overfit and the Out-of-Sample Error -- Strategy 1: Train-Test Split -- Strategy 2: Train-Validation-Test Split -- Strategy 3: Cross-Validation for Forecasting -- K-Fold Cross-Validation -- Time Series Cross-Validation -- Rolling Time Series Cross-Validation -- Backtesting -- Which Strategy to Use for Safe Forecasts? -- Final Considerations on Model Evaluation -- Key Takeaways -- Part II: Univariate Time Series Models -- Chapter 3: The AR Model -- Autocorrelation: The Past Influences the Present -- Compute Autocorrelation in Earthquake Counts -- Positive and Negative Autocorrelation -- Stationarity and the ADF Test -- Differencing a Time Series -- Lags

in Autocorrelation -- Partial Autocorrelation -- How Many Lags to Include? -- AR Model Definition -- Estimating the AR Using Yule-Walker Equations -- The Yule-Walker Method -- Train-Test Evaluation and Tuning -- Key Takeaways -- Chapter 4: The MA Model -- The Model Definition -- Fitting the MA Model -- Stationarity -- Choosing Between an AR and an MA Model. Application of the MA Model -- Multistep Forecasting with Model Retraining -- Grid Search to Find the Best MA Order -- Key Takeaways -- Chapter 5: The ARMA Model -- The Idea Behind the ARMA Model -- The Mathematical Definition of the ARMA Model -- An Example: Predicting Sunspots Using ARMA -- Fitting an ARMA(1,1) Model -- More Model Evaluation KPIs -- Automated Hyperparameter Tuning -- Grid Search: Tuning for Predictive Performance -- Key Takeaways -- Chapter 6: The ARIMA Model -- ARIMA Model Definition -- Model Definition -- ARIMA on the CO2 Example -- Key Takeaways -- Chapter 7: The SARIMA Model -- Univariate Time Series Model Breakdown -- The SARIMA Model Definition -- Example: SARIMA on Walmart Sales -- Key Takeaways -- Part III: Multivariate Time Series Models -- Chapter 8: The SARIMAX Model -- Time Series Building Blocks -- Model Definition -- Supervised Models vs. SARIMAX -- Example of SARIMAX on the Walmart Dataset -- Key Takeaways -- Chapter 9: The VAR Model -- The Model Definition -- Order: Only One Hyperparameter -- Stationarity -- Estimation of the VAR Coefficients -- One Multivariate Model vs. Multiple Univariate Models -- An Example: VAR for Forecasting Walmart Sales -- Key Takeaways -- Chapter 10: The VARMAX Model -- Model Definition -- Multiple Time Series with Exogenous Variables -- Key Takeaways -- Part IV: Supervised Machine Learning Models -- Chapter 11: The Linear Regression -- The Idea Behind Linear Regression -- Model Definition -- Example: Linear Model to Forecast CO2 Levels -- Key Takeaways -- Chapter 12: The Decision Tree Model -- Mathematics -- Splitting -- Pruning and Reducing Complexity -- Example -- Key Takeaways -- Chapter 13: The kNN Model -- Intuitive Explanation -- Mathematical Definition of Nearest Neighbors -- Combining k Neighbors into One Forecast -- Deciding on the Number of Neighbors k. Predicting Traffic Using kNN -- Grid Search on kNN -- Random Search: An Alternative to Grid Search -- Key Takeaways -- Chapter 14: The Random Forest -- Intuitive Idea Behind Random Forests -- Random Forest Concept 1: Ensemble Learning -- Bagging Concept 1: Bootstrap -- Bagging Concept 2: Aggregation -- Random Forest Concept 2: Variable Subsets -- Predicting Sunspots Using a Random Forest -- Grid Search on the Two Main Hyperparameters of the Random Forest -- Random Search CV Using Distributions -- Distribution for max_features -- Distribution for n_estimators -- Fitting the RandomizedSearchCV -- Interpretation of Random Forests: Feature Importance -- Key Takeaways -- Chapter 15: Gradient Boosting with XGBoost and LightGBM -- Boosting: A Different Way of Ensemble Learning -- The Gradient in Gradient Boosting -- Gradient Boosting Algorithms -- The Difference Between XGBoost and LightGBM -- Forecasting Traffic Volume with XGBoost -- Forecasting Traffic Volume with LightGBM -- Hyperparameter Tuning Using Bayesian Optimization -- The Theory of Bayesian Optimization -- Bayesian Optimization Using scikit-optimize -- Conclusion -- Key Takeaways -- Part V: Advanced Machine and Deep Learning Models -- Chapter 16: Neural Networks -- Fully Connected Neural Networks -- Activation Functions -- The Weights: Backpropagation -- Optimizers -- Learning Rate of the Optimizer -- Hyperparameters at Play in Developing a NN -- Introducing the Example Data -- Specific Data Prep Needs for a NN -- Scaling

and Standardization -- Principal Component Analysis (PCA) -- The Neural Network Using Keras -- Conclusion -- Key Takeaways -- Chapter 17: RNNs Using SimpleRNN and GRU -- What Are RNNs: Architecture -- Inside the SimpleRNN Unit -- The Example -- Predicting a Sequence Rather Than a Value -- Univariate Model Rather Than Multivariable -- Preparing the Data -- A Simple SimpleRNN. SimpleRNN with Hidden Layers -- Simple GRU -- GRU with Hidden Layers -- Key Takeaways -- Chapter 18: LSTM RNNs -- What Is LSTM -- The LSTM Cell -- Example -- LSTM with One Layer of 8 -- LSTM with Three Layers of 64 -- Conclusion -- Key Takeaways -- Chapter 19: The Prophet Model -- The Example -- The Prophet Data Format -- The Basic Prophet Model -- Adding Monthly Seasonality to Prophet -- Adding Holiday Data to Basic Prophet -- Adding an Extra Regressor to Prophet -- Tuning Hyperparameters Using Grid Search -- Key Takeaways -- Chapter 20: The DeepAR Model -- About DeepAR -- Model Training with DeepAR -- Predictions with DeepAR -- Probability Predictions with DeepAR -- Adding Extra Regressors to DeepAR -- Hyperparameters of the DeepAR -- Benchmark and Conclusion -- Key Takeaways -- Chapter 21: Model Selection -- Model Selection Based on Metrics -- Model Structure and Inputs -- One-Step Forecasts vs. Multistep Forecasts -- Model Complexity vs. Gain -- Model Complexity vs. Interpretability -- Model Stability and Variation -- Conclusion -- Key Takeaways -- Index.

Sommario/riassunto

Cover all the machine learning techniques relevant for forecasting problems, ranging from univariate and multivariate time series to supervised learning, to state-of-the-art deep forecasting models such as LSTMs, recurrent neural networks, Facebook's open-source Prophet model, and Amazon's DeepAR model. Rather than focus on a specific set of models, this book presents an exhaustive overview of all the techniques relevant to practitioners of forecasting. It begins by explaining the different categories of models that are relevant for forecasting in a high-level language. Next, it covers univariate and multivariate time series models followed by advanced machine learning and deep learning models. It concludes with reflections on model selection such as benchmark scores vs. understandability of models vs. compute time, and automated retraining and updating of models. Each of the models presented in this book is covered in depth, with an intuitive simple explanation of the model, a mathematical transcription of the idea, and Python code that applies the model to an example data set. Reading this book will add a competitive edge to your current forecasting skillset. The book is also adapted to those who have recently started working on forecasting tasks and are looking for an exhaustive book that allows them to start with traditional models and gradually move into more and more advanced models. What You Will Learn Carry out forecasting with Python Mathematically and intuitively understand traditional forecasting models and state-of-the-art machine learning techniques Gain the basics of forecasting and machine learning, including evaluation of models, cross-validation, and back testing Select the right model for the right use case Who This Book Is For The advanced nature of the later chapters makes the book relevant for applied experts working in the domain of forecasting, as the models covered have been published only recently. Experts working in the domain will want to update their skills as traditional models are regularly being outperformed by newer models.
