

Table of Contents

Preface	1
Chapter 1: Introduction to Time Series Analysis and R	7
Technical requirements	7
Time series data	8
Historical background of time series analysis	10
Time series analysis	11
Learning with real-life examples	12
Getting started with R	14
Installing R	15
A brief introduction to R	16
R operators	16
Assignment operators	16
Arithmetic operators	17
Logical operators	18
Relational operators	19
The R package	20
Installation and maintenance of a package	21
Loading a package in the R working environment	22
The key packages	23
Variables	24
Importing and loading data to R	25
Flat files	26
Web API	27
R datasets	28
Working and manipulating data	29
Querying the data	31
Help and additional resources	34
Summary	35
Chapter 2: Working with Date and Time Objects	37
Technical requirements	37
The date and time formats	38
Date and time objects in R	40
Creating date and time objects	43
Importing date and time objects	44
Reformatting and converting date objects	44
Handling numeric date objects	48
Reformatting and conversion of time objects	50
Time zone setting	53
Creating a date or time index	54

Manipulation of date and time with the lubridate package	55
Reformatting date and time objects – the lubridate way	56
Utility functions for date and time objects	59
Summary	61
Chapter 3: The Time Series Object	63
Technical requirement	64
The Natural Gas Consumption dataset	64
The attributes of the ts class	65
Multivariate time series objects	68
Creating a ts object	70
Creating an mts object	75
Setting the series frequency	76
Data manipulation of ts objects	80
The window function	80
Aggregating ts objects	82
Creating lags and leads for ts objects	83
Visualizing ts and mts objects	84
The plot.ts function	85
The dygraphs package	87
The TSstudio package	90
Summary	92
Chapter 4: Working with zoo and xts Objects	93
Technical requirement	93
The zoo class	94
The zoo class attributes	95
The index of the zoo object	96
Working with date and time objects	97
Creating a zoo object	99
Working with multiple time series objects	103
The xts class	104
The xts class attributes	105
The xts functionality	105
The periodicity function	106
Manipulating the object index	106
Subsetting an xts object based on the index properties	107
Manipulating the zoo and xts objects	108
Merging time series objects	109
Rolling windows	111
Creating lags	113
Aggregating the zoo and xts objects	113
Plotting zoo and xts objects	114
The plot.zoo function	115
The plot.xts function	116
xts, zoo, or ts – which one to use?	122

Summary	123
Chapter 5: Decomposition of Time Series Data	125
Technical requirement	125
The moving average function	126
The rolling window structure	126
The average method	128
The MA attributes	129
The simple moving average	132
Two-sided MA	137
A simple MA versus a two-sided MA	140
The time series components	142
The cycle component	143
The trend component	145
The seasonal component	147
The seasonal component versus the cycle component	150
White noise	152
The irregular component	156
The additive versus the multiplicative model	156
Handling multiplicative series	158
The decomposition of time series	160
Classical seasonal decomposition	161
Seasonal adjustment	165
Summary	166
Chapter 6: Seasonality Analysis	167
Technical requirement	167
Seasonality types	168
Seasonal analysis with descriptive statistics	171
Summary statistics tables	171
Seasonal analysis with density plots	180
Structural tools for seasonal analysis	185
Seasonal analysis with the forecast package	185
Seasonal analysis with the TSstudio package	187
Summary	195
Chapter 7: Correlation Analysis	197
Technical requirement	197
Correlation between two variables	198
Lags analysis	200
The autocorrelation function	202
The partial autocorrelation function	205
Lag plots	206
Causality analysis	211
Causality versus correlation	211

The cross-correlation function	213
Summary	218
Chapter 8: Forecasting Strategies	219
Technical requirement	219
The forecasting workflow	220
Training approaches	222
Training with single training and testing partitions	222
Forecasting with backtesting	225
Forecast evaluation	229
Residual analysis	229
Scoring the forecast	232
Forecast benchmark	235
Finalizing the forecast	238
Handling forecast uncertainty	240
Confidence interval	240
Simulation	241
Horse race approach	243
Summary	245
Chapter 9: Forecasting with Linear Regression	247
Technical requirement	247
The linear regression	248
Coefficients estimation with the OLS method	250
The OLS assumptions	252
Forecasting with linear regression	253
Forecasting the trend and seasonal components	253
Features engineering of the series components	255
Modeling the series trend and seasonal components	259
The tslm function	270
Modeling single events and non-seasonal events	271
Forecasting a series with multiseasonality components – a case study	273
The UKgrid series	274
Preprocessing and feature engineering of the UKdaily series	276
Training and testing the forecasting model	280
Model selection	284
Residuals analysis	286
Finalizing the forecast	287
Summary	289
Chapter 10: Forecasting with Exponential Smoothing Models	291
Technical requirement	291
Forecasting with moving average models	292
The simple moving average	292
Weighted moving average	299

Forecasting with exponential smoothing	301
Simple exponential smoothing model	301
Forecasting with the ses function	306
Model optimization with grid search	309
Holt method	311
Forecasting with the holt function	313
Holt-Winters model	318
Summary	328
Chapter 11: Forecasting with ARIMA Models	329
Technical requirement	330
The stationary process	330
Transforming a non-stationary series into a stationary series	334
Differencing time series	335
Log transformation	337
The random walk process	338
The AR process	341
Identifying the AR process and its characteristics	344
The moving average process	346
Identifying the MA process and its characteristics	348
The ARMA model	349
Identifying an ARMA process	352
Manual tuning of the ARMA model	353
Forecasting AR, MA, and ARMA models	357
The ARIMA model	358
Identifying an ARIMA process	359
Identifying the model degree of differencing	359
The seasonal ARIMA model	364
Tuning the SARIMA model	365
Tuning the non-seasonal parameters	365
Tuning the seasonal parameters	365
Forecasting US monthly natural gas consumption with the SARIMA model – a case study	366
The auto.arima function	377
Linear regression with ARIMA errors	379
Violation of white noise assumption	379
Modeling the residuals with the ARIMA model	382
Summary	387
Chapter 12: Forecasting with Machine Learning Models	389
Technical requirement	389
Why and when should we use machine learning?	390
Why h2o?	391
Forecasting monthly vehicle sales in the US – a case study	392
Exploratory analysis of the USVSAles series	392
The series structure	392

The series components	394
Seasonal analysis	395
Correlation analysis	396
Exploratory analysis – key findings	397
Feature engineering	399
Training, testing, and model evaluation	401
Model benchmark	402
Starting a h2o cluster	403
Training an ML model	405
Forecasting with the Random Forest model	406
Forecasting with the GBM model	413
Forecasting with the AutoML model	416
Selecting the final model	419
Summary	421
Other Books You May Enjoy	423
Index	427