

Data Science for Business Leaders

Lecture Notes

Prof. Dr. Stephan Huber

June 18, 2024

Table of contents

1. (Extended) Syllabus	5
1.1. Syllabus	5
1.2. How I let ChatGPT wrote the extended syllabus	6
1.3. The ChatGPT generated extended syllabus	7
1.3.1. Scope and Nature of Data Science	7
1.3.2. Emerging Trends in a Data-Driven Business Environment	9
1.3.3. Data Science Process in Business	11
1.3.4. Data Literacy	13
1.3.5. Overview of Data Science Methods	15
1.3.6. Introduction to Data Scientific Tools	16
1.4. How to use Natural Language Processing tools	18
2. Data Pitfalls	20
2.1. Epistemic Errors	20
2.2. Technical trespasses	22
2.3. Mathematical Miscues	25
2.4. Statistical slipups	26
2.4.1. Descriptive debacles	26
2.4.2. Inferential infernos	26
2.4.3. Slippery sampling	27
2.4.4. Insensitivity to sample size	27
2.5. Analytical Aberrations	28
2.5.1. The intuition/analysis false dichotomy	28
2.5.2. Exuberant extrapolations	28
2.5.3. Ill-advised interpolations	29
2.5.4. Funky forecasts	29
2.5.5. Moronic measures	29
2.6. Graphical Gaffes	29
2.6.1. Challenging charts	29
2.6.2. Data dogmatism	30
2.6.3. The optimize/satisfice false dichotomy	30
2.7. Design Dangers	30
2.7.1. Confusing Colors	30
2.7.2. Usability uh-ohs	31
3. How to use R for data science	32
4. Collaborating with Git and GitHub	33
4.1. Introduction	33
4.2. Install Git	34
4.3. Using Git from the terminal	35
4.3.1. Configuring Git	35
4.3.2. Initializing a Repository	36
4.3.3. Staging Changes	36

Table of contents

4.3.4. Committing Changes	37
4.3.5. Pushing Changes	37
4.3.6. Undo changes	37
4.4. Using Git from RStudio	38
4.4.1. Set up Git in RStudio	38
4.4.2. Connecting RStudio Projects with GitHub repositories	38
4.5. Make a contribution using Git and GitHub	39
5. Markdown, Quarto, and R Markdown	42
5.1. Why Markdown and Quarto?	42
5.1.1. No-code vs. code-based writing applications	42
5.1.2. Typical (mis)usage of WYSIWYG applications	43
5.1.3. Advantages of Quarto for writing text	43
5.2. Markdown	44
5.3. Quarto	44
5.3.1. Introduction	44
5.3.2. Create an APA compliant manuscript using Quarto	45
5.4. R Markdown	46
6. Create and host a website	51
6.1. Creating a website with Quarto	51
6.2. Hosting the website on GitHub	51
References	53
Appendices	54
A. Exam(ple) question	54
A.0.1. Question	54
A.0.2. Solutions	54


List of figures

1.	Prof. Dr. Stephan Huber	2
2.1.	Bananas in various stages of ripeness	22
2.2.	Distribution of bullet holes in returned aircraft	24
4.1.	The FINAL.doc problem	33
4.2.	GitHub is big	33
4.3.	Memorizing six git commands	34
4.4.	Three git commands you really need	35
4.5.	Copy the https URL of your repo	37
4.6.	Fork the repo	39
5.1.	Example of an R Markdown file	46
5.2.	R Markdown Cheatsheet from Posit	46
5.3.	Xie et al. [2020]: R Markdown Cookbook	47
5.4.	Xie et al. [2018]: R Markdown: The Definitive Guide	47

List of tables

4.1. Most important git commands	35
4.2. Most common bash commands	36

Preface



About Data Science for Business Leaders

This course provides an overview of the tools and methods that data science offers to business leaders. Of course, we only scratch the surface on many topics. Many topics require more in-depth analysis to fully grasp the underlying principles and their benefits. In my opinion, business leaders often don't understand every detail of the decisions they make. They can't do that and often they don't need to. However, they are good at identifying and prioritizing the key factors that can drive success. They have a sense of the essential elements that have a significant impact. Ultimately, leaders lead. They recognize, cultivate and manage the expertise of their employees and teammates. They orchestrate the knowledge around them, they lead their teams effectively and help them avoid critical missteps.

About the cover and the logo of the notes

I realize that having a visually appealing cover and logo does not directly help you become an exceptional leader. The images themselves do not teach you how to master the tools of data science or apply its methods effectively. However, I do believe that attractive visual elements appeal to many students and individuals. Therefore, as a teacher, I think it is wise to occasionally incorporate these elements into my notes. This approach may encourage you to continue reading and studying the material. By the way, it only took me a couple of minutes to create these two images using ChatGPT.

About the notes

A PDF version of these notes is available [here](#)..

Please note that while the PDF contains the same content, it has not been optimized for PDF format. Therefore, some parts may not appear as intended.

- These notes aims to support my lecture at the HS Fresenius but are incomplete and no substitute for taking actively part in class.
- I hope you find this book helpful. Any feedback is both welcome and appreciated.
- This is work in progress so please check for updates regularly.
- These notes offer a curated collection of explanations, exercises, and tips to facilitate learning R without causing unnecessary frustration. However, these notes don't aim to rival comprehensive textbooks.
- These notes are published under the [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#). This means it can be reused, remixed, retained, revised and redistributed as long as appropriate credit is given to the authors. If you remix, or modify the original version of this open textbook, you must redistribute all

versions of this open textbook under the same license.



- I host the notes in a [GitHub repo](#).

Structure of these notes

About the author

i Contact:

Prof. Dr. Stephan Huber
Hochschule Fresenius für Wirtschaft & Medien GmbH
Im MediaPark 4c
50670 Cologne
Office: 4e OG-3
Telefon: +49 221 973199-523
Mail: stephan.huber@hs-fresenius.de
Private homepage: www.hubchev.github.io
Github: <https://github.com/hubchev>

Figure 1.: Prof. Dr. Stephan Huber



I am a Professor of *International Economics and Data Science* at HS Fresenius, holding a Diploma in Economics from the University of Regensburg and a Doctoral Degree (summa cum laude) from the University of Trier. I completed postgraduate studies at the Interdisciplinary Graduate Center of Excellence at the Institute for Labor Law and Industrial Relations in the

Preface

European Union (IAAEU) in Trier. Prior to my current position, I worked as a research assistant to Prof. Dr. Dr. h.c. Joachim Möller at the University of Regensburg, a post-doc at the Leibniz Institute for East and Southeast European Studies (IOS) in Regensburg, and a freelancer at Charles University in Prague.

Throughout my career, I have also worked as a lecturer at various institutions, including the TU Munich, the University of Regensburg, Saarland University, and the Universities of Applied Sciences in Frankfurt and Augsburg. Additionally, I have had the opportunity to teach abroad for the University of Cordoba in Spain, the University of Perugia in Italy, and the Petra Christian University in Surabaya, Indonesia. My published work can be found in international journals such as the Canadian Journal of Economics and the Stata Journal. For more information on my work, please visit my private homepage at hubchev.github.io.

I was always fascinated by data and statistics. For example, in 1992 I could name all soccer players in Germany's first division including how many goals they scored. Later, in 2003 I joined the introductory statistics course of [Daniel Rösch](#). I learned among others that probabilities often play a role when analyzing data. I continued my data science journey with [Harry Haupt's](#) *Introductory Econometrics* course, where I studied the infamous Jeffrey M. [Wooldridge \[2002\]](#) textbook. It got me hooked and so I took all the courses [Rolf Tschernig](#) offered at his chair of Econometrics, where I became a tutor at the University of Regensburg and a research assistant of [Joachim Möller](#). Despite everything we did had to do with how to make sense out of data, we never actually used the term *data science* which is also absent in the more 850 pages long textbook by [Wooldridge \[2002\]](#). The book also remains silent about *machine learning* or *artificial intelligence*. These terms became popular only after I graduated. The *Harvard Business Review* article by [Davenport and Patil \[2012\]](#) who claimed that data scientist is “The Sexiest Job of the 21st Century” may have boosted the popularity.

The term “data scientist” has become remarkably popular, and many people are eager to adopt this title. Although I am a professor of data science, my professional identity is more like that of an applied, empirically-oriented international economist. My hesitation to adopt the title “data scientist” also stems from the deep respect I have developed through my interactions with econometricians and statisticians. Considering their in-depth expertise, I feel like a passionate amateur.

Ultimately, I poke around in data to find something interesting. Much like my ten-year-old younger self who analyzed soccer statistics to gain a deeper understanding of the sport. The only thing that has changed since then is that I know more promising methods and can efficiently use tools for data processing and data analysis.

To students

I'm not a business leader myself. I'm a professor, and like most of my peers, I don't run a big company nor do I hold a top position in a multinational corporation. Actually, I believe that leading a company successfully full-time is not compatible with being a full-time professor, and vice versa. Despite this, in my role as a professor and study program director, I do have certain responsibilities and occasionally lead people, particularly students. This, however, doesn't let me feel like I am a business leader. Is this necessary to teach you a course entitled “Data Science for Business Leaders”? I don't think so: Firstly, if you want to become a business leader, do not pick a role model that works in academics like I do. Secondly, I believe I can provide you with valuable insights into data science to support you on the journey of becoming a business leader. And thirdly, this course is not entitled “I am a business leader who teaches you data science”. Thus, this course is “for” business leaders and “for” you and not about me or my personality.

Preface

Enjoy.

1. (Extended) Syllabus

Some time ago, I wrote the syllabus of this course, see Section 1.1. that was successfully accredited by the state. Since the syllabus is quite short, I thought about an extended version to help students familiarize themselves with the buzzwords I used in the original version. As a data scientist who is supposed to teach you how to use the tools of data science in business, I wondered how I could write it without devoting too many resources to it, that is, being economical. I thought it would be a good idea and a fun exercise to let ChatGPT from OpenAI [2024] write it for me. ChatGPT uses a large language model to generate text based on a vast corpus of text data. I provided ChatGPT with specific prompts to craft the extended abstract, which you can find in Section 1.2. The final extended syllabus, actually written by ChatGPT, is available in Section 1.3.

1.1. Syllabus

Scope and Nature of Data Science

- Defining data science as an academic discipline (informatics, computer science, mathematics, statistics, econometrics, social science)
- Importance of data science in businesses

Emerging Trends in a Data-Driven Business Environment

- Evolution of computers, computing, and data processing
- Business intelligence (performance marketing, etc.)
- Artificial intelligence, machine learning, deep learning, and algorithms
- Big data
- Internet of things, cloud computing, blockchain
- Industry 4.0 and remote working

Data Science Process in Business

- Workflows and data science life cycles (OSEMN, CRISP-DM, Kanban, TDSP, ...)
- Types of data science roles (data engineer, data analyst, machine learning engineer, business intelligence analyst, database administrator, data product manager, ...)

Data Literacy

- Conceptual framework (knowledge and understanding of data and applications of data)
- Data collection (identify, collect, and assess data)
- Data management (organize, clean, convert, curate, and preserve data)
- Data evaluation (plan, conduct, evaluate, and assess data analyses)
- Data application (share, reflect, and evaluate results of analyses and compare them with other findings considering ethical issues and scientific standards)

Overview of Data Science Methods

1. (Extended) Syllabus

- Data exploration and data mining
- Supervised and unsupervised learning
- Regression and classification
- Predictive analysis
- Causal analysis

Introduction to Data Scientific Tools

- Writing and publishing reports (Markdown, Quarto)
- Collaborating in teams using a version control system (git)
- Overview of programming languages (R, Python, SQL, ...)
- Overview of no-code and low-code tools for data science (makeML, PyCaret, Rapidminer, KNIME, etc.)
- Development environments (Unix-like systems, containers, APIs, Jupyter, Rstudio, etc.)

1.2. How I let ChatGPT wrote the extended syllabus

Here are the seven prompts to ChatGPT 4.0:

1. Act as a professor of data science. Write lecture notes for management students. The first chapter of the notes should contain the following:

Scope and nature of data science

- Defining data science as an academic discipline (informatics, computer science, mathematics, statistics, econometrics, social science)
- Importance of data science in businesses

2. The chapter of the lecture notes should contain the following:

Emerging Trends in a Data-Driven Business Environment

- Evolution of computers, computing, and data processing
- Business Intelligence (Performance Marketing, etc.)
- Artificial intelligence, machine learning, deep learning, and algorithms
- Big data
- Internet of things, cloud computing, blockchain
- Industry 4.0 and remote working

3. The third chapter of the lecture notes should contain the following:

Data science process in business

- Workflows and Data science life cycles (OSEMN, CRISP-DM, Kanban, TDSP, ...)
- Types of data science roles (data engineer, data analyst, machine learning engineer, business intelligence analyst, database administrator, data product manager, ...)

4. Do the same for the fourth section, which contains:

Data literacy

- Conceptual framework (knowledge and understanding of data and applications of data)
- Data collection (identify, collect, and assess data)
- Data management (organize, clean, convert, curate, and preserve Kdata)

1. (Extended) Syllabus

- Data evaluation (plan, conduct, evaluate, and assess data analyses)
- Data application (share, reflect, and evaluate results of analyses and compare them with other findings considering ethical issues and scientific standards)

5. And the next section:

Overview of data science methods

- Data exploration and data mining
- Supervised and unsupervised learning
- Regression and classification
- Predictive analysis
- Causal analysis

6. And the final section:

Introduction to data scientific tools

- Writing and publishing reports (Markdown, Quarto)
- Collaborating in teams using a version control system (git)
- Overview on Programming languages (R, Python, SQL, ...)
- Overview on no-code and low-code tools for data science (makeML, PyCaret, Rapidminer, KNIME, etc.)
- Development environments (Unix-like systems, containers, APIs, Jupyter, Rstudio, etc.)

7. I am not happy with the third section. In particular the section “Workflows and Data Science Life Cycles”. Can you re-write it? Consider the learning objective a bit better, that is, summarize the strengths, weaknesses and use-cases of data science processes and workflows and demonstrate the ability to justify and explain tools and methods used along the process.

I manually cut some parts of the last answer and pasted it into the respective section.

I assembled everything by copy and paste it step-by-step into Quarto and render it. That was all I did to create the notes.

Exercise 1.1. What do think about the extended syllabus and my approach
What do you think about the extended syllabus? Is it well-written? Does it include a lot of details or errors? What aspects do you think we should work on manually? How could my prompts be improved to yield better results?

1.3. The ChatGPT generated extended syllabus

The following text was generated as described in Section 1.2 using [OpenAI \[2024\]](#).

1.3.1. Scope and Nature of Data Science

Welcome to the introductory chapter on Data Science, designed specifically for management students. In this chapter, we will explore the multifaceted discipline of data science, understanding its definitions, scope, and the pivotal role it plays in the business world today.

1. (Extended) Syllabus

1.3.1.1. Defining Data Science

Data Science is an interdisciplinary field that uses scientific methods, processes, algorithms, and systems to extract knowledge and insights from structured and unstructured data. It is a confluence of several disciplines including informatics, computer science, mathematics, statistics, econometrics, and social science. This integration allows for a comprehensive approach to solving complex problems, making informed decisions, and generating actionable insights through data analysis.

1.3.1.1.1. Informatics

It is the study of computational systems and the way humans interact with these systems. In data science, informatics plays a critical role in the management and processing of data.

1.3.1.1.2. Computer Science

This foundational pillar focuses on the development of algorithms and software that are used to process, store, and analyze data.

1.3.1.1.3. Mathematics and Statistics

Mathematics provides the theoretical foundation for algorithms and models, while statistics is crucial for understanding and interpreting data. Together, they enable data scientists to perform quantitative analysis and predictive modeling.

1.3.1.1.4. Econometrics

Econometrics applies statistical methods to economic data, enabling the analysis of economic relationships. In data science, econometrics helps in forecasting economic trends and evaluating economic policies through data.

1.3.1.1.5. Social Science

The incorporation of social science principles allows data scientists to understand human behavior and societal trends, which is invaluable for businesses looking to market their products or services effectively.

1.3.1.2. Importance of Data Science in Businesses

Data Science has become a cornerstone in the business world, offering unprecedented opportunities for growth, efficiency, and innovation. Here's how:

1.3.1.2.1. Data-Driven Decision Making

Businesses can make more informed decisions by analyzing vast amounts of data. This leads to better strategic planning, operational efficiency, and competitive advantage.

1. (Extended) Syllabus

1.3.1.2.2. Customer Insights

Through data analysis, companies can gain deep insights into customer behavior, preferences, and trends. This information is critical for tailoring marketing strategies, improving customer experiences, and boosting customer loyalty.

1.3.1.2.3. Risk Management

Data science enables businesses to identify and mitigate risks by analyzing historical data, detecting patterns, and predicting future outcomes. This is particularly useful in finance, insurance, and cybersecurity.

1.3.1.2.4. Innovation and Product Development

By understanding market needs and customer feedback through data analysis, businesses can innovate and develop new products and services that meet or exceed customer expectations.

1.3.1.2.5. Operational Efficiency

Data science can optimize business operations, reduce costs, and improve service delivery through the analysis of operational data, leading to streamlined processes and better resource management.

1.3.1.2.6. Competitive Edge

In a data-driven world, businesses that leverage data science effectively can stay ahead of the curve, identifying trends and opportunities before their competitors.

In conclusion, data science is an essential academic discipline that integrates various fields to analyze and interpret complex data. Its importance in the business sector cannot be overstated, as it enables organizations to make informed decisions, understand their customers, manage risks, foster innovation, improve operations, and maintain a competitive edge in the market. As we delve deeper into this course, we will explore the methodologies, tools, and applications of data science in greater detail, equipping you with the knowledge and skills to leverage data science in your future management roles.

1.3.2. Emerging Trends in a Data-Driven Business Environment

In this chapter, we explore the dynamic landscape of a data-driven business environment. The rapid evolution of technology and its integration into business processes have led to the emergence of new trends that are shaping the future of industries across the globe. Let's dive into these trends and understand their significance.

1.3.2.1. Evolution of Computers, Computing, and Data Processing

The journey from the first generation of computers to the modern era of quantum computing marks a significant evolution in technology. Initially, computers were large, expensive, and limited in their capabilities. Over the decades, advancements in semiconductor technology, the invention of the microprocessor, and the development of personal computers transformed computing into an accessible and essential tool for businesses. Today, cloud computing and edge computing have further revolutionized data processing, allowing for more efficient data storage, access, and analysis.

1.3.2.2. Business Intelligence

Business Intelligence (BI) refers to the use of data analysis in business to support decision-making processes. BI tools analyze historical and current data to provide actionable insights, helping businesses to improve their performance. Performance marketing, a subset of BI, focuses on analyzing marketing campaigns in real time to optimize marketing strategies and expenditures for better ROI.

1.3.2.3. Artificial Intelligence, Machine Learning, Deep Learning, and Algorithms

Artificial Intelligence (AI) and its subsets, Machine Learning (ML) and Deep Learning (DL), are at the forefront of technological innovation. AI involves creating systems capable of performing tasks that typically require human intelligence. ML and DL are about teaching computers to learn from data, improving their accuracy over time without being explicitly programmed. These technologies are transforming business operations, from customer service automation and predictive analytics to personalized marketing and decision-making processes.

1.3.2.4. Big Data

Big Data refers to the vast volumes of data generated every second from various sources like social media, business transactions, and IoT devices. The ability to process and analyze this data has unlocked new opportunities for businesses to gain insights into customer behavior, market trends, and operational efficiency. Big data analytics is now a crucial tool for strategic planning and competitive analysis.

1.3.2.5. Internet of Things, Cloud Computing, Blockchain

- **Internet of Things (IoT):** IoT technology connects everyday devices to the internet, enabling them to send and receive data. This interconnectivity offers businesses real-time insights into their operations, asset tracking, and supply chain management.
- **Cloud Computing:** Cloud computing provides businesses with scalable computing resources over the internet, facilitating remote data storage, processing, and management. It supports flexibility, reduces IT costs, and enhances collaboration.
- **Blockchain:** Blockchain technology offers a secure, decentralized platform for transactions. It's particularly valuable for enhancing transparency, security, and efficiency in business operations, supply chain management, and financial services.

1.3.2.6. Industry 4.0 and Remote Working

Industry 4.0, also known as the fourth industrial revolution, integrates digital technologies into manufacturing and industry, including IoT, AI, and robotics. It represents a shift towards smart, automated production processes and data exchange. Coupled with the rise of remote working, Industry 4.0 technologies enable businesses to operate more flexibly, with teams collaborating effectively from various locations, leveraging digital tools and platforms for communication and project management.

In conclusion, the evolution of technology and its integration into business practices have brought about significant changes in the way companies operate. From the way data is processed and analyzed to the automation of manufacturing processes and the flexibility of remote working, these emerging trends are shaping the future of a data-driven business environment. As we move forward, understanding and adapting to these trends will be crucial for businesses looking to thrive in the digital age.

1.3.3. Data Science Process in Business

This chapter delves into the structured approach behind the application of data science in business settings. We will explore various data science workflows and life cycles that guide the process from raw data to actionable insights. Additionally, we will outline the different roles within a data science team and their contributions to this process.

1.3.3.1. Workflows and Data Science Life Cycles

Data science projects in business environments follow structured workflows and life cycles to ensure that the analysis is efficient, reproducible, and scalable. Several frameworks guide these processes, each with its strengths and applications.

1.3.3.1.1. OSEMN Framework

OSEMN (Obtain, Scrub, Explore, Model, iNterpret) is a streamlined approach to data science projects:

1. **Obtain:** Acquiring the data from various sources.
 2. **Scrub:** Cleaning the data to ensure it is accurate and usable.
 3. **Explore:** Analyzing the data to find patterns and relationships.
 4. **Model:** Applying statistical models to predict or classify data.
 5. **Interpret:** Drawing conclusions and making recommendations based on the model's results.
- **Strengths:** The OSEMN (Obtain, Scrub, Explore, Model, iNterpret) framework is straightforward and easy to understand, making it accessible for teams of all skill levels. It covers the essential steps of a data science project in a logical sequence.
 - **Weaknesses:** Its simplicity may overlook the complexity of certain stages, such as model validation or deployment.
 - **Use-Cases:** Ideal for small to medium-sized projects where the primary goal is to gain insights from data through exploration and modeling.

1.3.3.1.2. CRISP-DM

CRISP-DM stands for Cross-Industry Standard Process for Data Mining. It's a comprehensive framework that includes six phases:

1. **Business Understanding:** Define the project objectives and requirements.
 2. **Data Understanding:** Collect and explore the data.
 3. **Data Preparation:** Clean and preprocess the data.
 4. **Modeling:** Select and apply modeling techniques.
 5. **Evaluation:** Assess the model's performance.
 6. **Deployment:** Implement the model in a real-world setting.
- **Strengths:** CRISP-DM (Cross-Industry Standard Process for Data Mining) is industry-agnostic and provides a detailed structure that includes understanding the business problem and deploying the solution. It encourages iterative learning and refinement.
 - **Weaknesses:** Can be perceived as too rigid for projects requiring rapid development and deployment. The model doesn't explicitly address the updating or maintenance of deployed solutions.
 - **Use-Cases:** Suitable for projects that require close alignment with business objectives and thorough consideration of deployment strategies.

1.3.3.1.3. Kanban

Kanban is a lean method to manage and improve work across human systems. In data science, it helps in visualizing work, limiting work-in-progress, and maximizing efficiency.

- **Strengths:** Kanban is highly flexible and promotes continuous delivery. It allows teams to adapt quickly to changes and prioritize tasks effectively.
- **Weaknesses:** Without strict stages or phases, projects might lack direction or oversight, potentially leading to inefficiencies.
- **Use-Cases:** Best for dynamic environments where priorities shift frequently and teams must remain agile to respond to business needs.

1.3.3.1.4. TDSP (Team Data Science Process)

TDSP is a standardized approach to data science projects that helps teams to improve quality and efficiency. It includes:

- **Strengths:** TDSP offers a structured approach with a strong emphasis on standardized documentation and project management methodologies, facilitating collaboration and scalability.
- **Weaknesses:** Its comprehensive nature might introduce overhead and slow down smaller projects.
- **Use-Cases:** Ideal for larger teams working on complex projects that require coordination across different roles and departments.

1.3.3.2. Types of Data Science Roles

In a business environment, a data science team might consist of various specialized roles, each contributing uniquely to the data science process.

1. (Extended) Syllabus

1.3.3.2.1. Data Engineer

Focuses on the design, construction, and maintenance of the systems that data analysts and data scientists use for their work. They ensure that data flows smoothly from source to database to analytics.

1.3.3.2.2. Data Analyst

Works on processing and performing statistical analysis on existing datasets. They interpret the data to help the business make more informed decisions.

1.3.3.2.3. Machine Learning Engineer

Develops algorithms and predictive models to solve specific business problems using machine learning techniques.

1.3.3.2.4. Business Intelligence Analyst

Analyzes data to provide insights that help businesses with strategic planning. They use BI tools to convert data into understandable reports and dashboards.

1.3.3.2.5. Database Administrator

Responsible for managing, backing up, and ensuring the availability of the data stored in an organization's databases.

1.3.3.2.6. Data Product Manager

Oversees the development of data-driven products or services, ensuring that they meet the users' needs and the business objectives.

In summary, the data science process in business involves a structured approach to turning data into actionable insights. This process is supported by various frameworks and relies on the collaboration of professionals in specialized roles. Understanding these aspects of data science is crucial for anyone looking to leverage this discipline in a business context.

1.3.4. Data Literacy

Data literacy is the ability to read, understand, create, and communicate data as information. It encompasses a broad range of skills necessary for effectively working with data, from the initial stages of data collection to the final stages of analyzing and sharing findings. In this chapter, we will break down the conceptual framework of data literacy and explore its various components in detail.

1. (Extended) Syllabus

1.3.4.1. Conceptual Framework

At the heart of data literacy is a deep knowledge and understanding of how data can be used to make decisions, solve problems, and communicate ideas. This conceptual framework involves:

- **Understanding the nature of data:** Recognizing different types of data (quantitative vs. qualitative) and their sources.
- **Comprehending the applications of data:** Knowing how data can be used in various contexts to derive insights and inform decisions.

1.3.4.2. Data Collection

The first step in the data lifecycle involves identifying, collecting, and assessing data:

- **Identify:** Determining the data needed to answer a question or solve a problem.
- **Collect:** Gathering data from various sources, whether they are existing datasets or new data collected through surveys, experiments, or observations.
- **Assess:** Evaluating the quality of the data, including its relevance, accuracy, and completeness.

1.3.4.3. Data Management

Once data is collected, it must be managed effectively:

- **Organize:** Arranging data in a structured format that facilitates analysis.
- **Clean:** Removing errors or inconsistencies in the data.
- **Convert:** Transforming data into a format suitable for analysis.
- **Curate:** Selecting, annotating, and maintaining valuable data for current and future use.
- **Preserve:** Ensuring that data remains accessible and usable over time.

1.3.4.4. Data Evaluation

Evaluation is critical to understanding what the data signifies:

- **Plan:** Designing a methodology for analyzing the data.
- **Conduct:** Performing the analysis using appropriate statistical methods and tools.
- **Evaluate:** Assessing the quality and reliability of the analysis.
- **Assess:** Interpreting the results in the context of the research question or business problem.

1.3.4.5. Data Application

The final step involves applying the insights gained from data analysis:

- **Share:** Communicating findings to stakeholders through reports, presentations, or visualizations.
- **Reflect:** Considering the implications of the results and how they can inform future actions.
- **Evaluate results:** Comparing findings with those from other studies or data analyses to draw broader conclusions.

1. (Extended) Syllabus

- **Ethical considerations:** Ensuring that the use of data respects privacy, confidentiality, and ethical standards.
- **Scientific standards:** Adhering to rigorous standards of validity, reliability, and objectivity in data handling and analysis.

In summary, data literacy is a comprehensive set of skills that enable individuals to navigate the complex world of data from collection to application. By understanding and applying the concepts outlined in this chapter, individuals can enhance their ability to make informed decisions, solve problems, and communicate effectively using data.

1.3.5. Overview of Data Science Methods

Data science encompasses a wide array of methods and techniques for analyzing data, drawing insights, and making predictions. This chapter provides an overview of some core data science methods, including data exploration, data mining, machine learning approaches, and various types of analyses.

1.3.5.1. Data Exploration and Data Mining

1.3.5.1.1. Data Exploration

Data exploration involves analyzing data sets to find initial patterns, characteristics, and points of interest without making any prior assumptions. It typically includes summarizing the main characteristics of the data through visualizations and statistics.

1.3.5.1.2. Data Mining

Data mining is the process of discovering patterns and knowledge from large amounts of data. The data sources can include databases, data warehouses, the internet, and other sources. Data mining techniques include clustering, classification, regression, and association rule learning.

1.3.5.2. Supervised and Unsupervised Learning

1.3.5.2.1. Supervised Learning

Supervised learning is a type of machine learning where the algorithm learns from labeled training data, helping to predict outcomes for unforeseen data. It is divided into two main categories: regression and classification.

1.3.5.2.2. Unsupervised Learning

Unsupervised learning involves training on data without labeled responses. The system tries to learn the patterns and the structure from the data without any supervision. Common unsupervised learning methods include clustering and dimensionality reduction.

1.3.5.3. Regression and Classification

1.3.5.3.1. Regression

Regression methods are used to predict a continuous outcome variable based on one or more predictor variables. The goal is to find the relationship between variables and forecast an outcome. Linear regression is one of the most basic types of regression analysis.

1.3.5.3.2. Classification

Classification methods are used to predict or identify the category to which a new observation belongs. Examples include spam detection in email service providers and customer churn prediction.

1.3.5.4. Predictive Analysis

Predictive analysis uses statistical algorithms and machine learning techniques to identify the likelihood of future outcomes based on historical data. It's used in various fields, including finance, healthcare, marketing, and weather forecasting, to make more informed decisions.

1.3.5.5. Causal Analysis

Causal analysis seeks to identify and understand the cause-and-effect relationships between variables. Unlike correlation, which merely indicates that two variables move together, causation establishes that a change in one variable is responsible for a change in another.

In conclusion, these methods and techniques form the backbone of data science, enabling professionals to extract valuable insights, make predictions, and inform decision-making processes. Understanding these methods is crucial for anyone looking to delve into data science or apply its principles in their field.

1.3.6. Introduction to Data Scientific Tools

The practice of data science requires not only a solid understanding of theories and methodologies but also proficiency in a variety of tools and technologies. This chapter introduces essential tools for writing and publishing reports, collaborating in teams, programming, as well as no-code and low-code platforms, and development environments.

1.3.6.1. Writing and Publishing Reports

1.3.6.1.1. Markdown

Markdown is a lightweight markup language with plain-text formatting syntax. Its simplicity and ease of conversion to HTML and other formats make it an ideal choice for writing and publishing reports, documentation, and articles.

1. (Extended) Syllabus

1.3.6.1.2. Quarto

Quarto is an open-source scientific and technical publishing system built on Pandoc. It enables users to create dynamic and reproducible reports and articles that can include executable code from various programming languages, such as R and Python.

1.3.6.2. Collaborating in Teams Using a Version Control System

1.3.6.2.1. Git

Git is a distributed version control system that enables multiple developers to work together on the same project efficiently. It tracks changes in source code during software development, supporting collaboration and fostering code integrity.

1.3.6.3. Overview of Programming Languages

1.3.6.3.1. R

R is a programming language and free software environment for statistical computing and graphics, widely used among statisticians and data miners.

1.3.6.3.2. Python

Python is a high-level, interpreted programming language known for its simplicity and versatility. It has a wide range of libraries for data analysis, machine learning, and data visualization, making it a popular choice in data science.

1.3.6.3.3. SQL

SQL (Structured Query Language) is the standard language for managing and manipulating relational databases. It allows users to query, update, and manage data.

1.3.6.4. Overview of No-Code and Low-Code Tools for Data Science

1.3.6.4.1. makeML

A no-code platform for machine learning, makeML simplifies the process of training and deploying ML models without writing extensive code.

1.3.6.4.2. PyCaret

PyCaret is a low-code machine learning library in Python that automates machine learning workflows. It enables data scientists to perform end-to-end experiments quickly and efficiently.

1.3.6.4.3. Rapidminer

Rapidminer is a data science platform that provides an integrated environment for data preparation, machine learning, deep learning, text mining, and predictive analytics.

1. (Extended) Syllabus

1.3.6.4.4. KNIME

KNIME is an open-source, graphical workbench for the entire analysis process: data access, data transformation, initial investigation, powerful predictive analytics, visualization, and reporting.

1.3.6.5. Development Environments

1.3.6.5.1. Unix-like Systems

Unix-like operating systems, including Linux and macOS, provide powerful tools and environments for software development and data science.

1.3.6.5.2. Containers

Containers, such as Docker, allow for the packaging of applications and their dependencies in a virtual container that can run on any Linux server, enabling easy deployment and scalability.

1.3.6.5.3. APIs

Application Programming Interfaces (APIs) enable different software applications to communicate with each other, facilitating data exchange and integration.

1.3.6.5.4. Jupyter

Jupyter Notebook is an open-source web application that allows you to create and share documents containing live code, equations, visualizations, and narrative text.

1.3.6.5.5. RStudio

RStudio is an integrated development environment (IDE) for R. It provides a user-friendly interface for coding, debugging, and visualizing data.

In summary, the array of tools and technologies available to data scientists is broad and varied, catering to different aspects of the data science workflow. From data manipulation and analysis to collaboration and report writing, mastering these tools is essential for effective data science practice.

1.4. How to use Natural Language Processing tools

AI text generation tools like ChatGPT can be valuable resources for students and business professionals. These tools provide quick access to information on a wide range of topics and can summarize long texts, articles, or papers, saving time. The interactive nature of these tools can help explain complex concepts in an easy-to-understand manner and generate ideas for essays, projects, and research topics. They are particularly helpful in creating structured outlines for assignments, papers, and meetings. Moreover, these tools are flexible, aiding in writing and improving initial drafts by suggesting enhancements in grammar, punctuation, and writing style.

1. (Extended) Syllabus

When used wisely, AI tools can save time on preliminary writing tasks and research, allowing students to focus on deeper analysis and learning. However, reliance on these tools requires careful consideration to avoid potential pitfalls. Here are some important points to consider:

Plagiarism: Submitting AI-generated text as original work can be considered plagiarism. Educational institutions value the development of students' own ideas and critical thinking skills. Overreliance on AI tools may hinder these skills. Universities and supervisors may restrict the use of AI tools; hence, students should be aware of their institution's policies. It's recommended to use AI tools as a supplement to one's research and writing, not as a replacement. Properly attribute any AI-generated content that contributes to your work to maintain academic honesty and integrity.

Accuracy: AI tools may produce incorrect or misleading information and may reflect biases present in the training data. Therefore, it is essential to cross-check and verify the accuracy of AI-generated content using credible sources.

Contextual relevance and personalization: AI may not fully understand the specific context or requirements of an assignment, project, paper, or thesis. Actively work on improving and personalizing the generated content to meet the required quality standards. Carefully review the AI-generated text for errors, coherence, and flow. In particular, consider the following:

- Organize the content logically with a clear introduction, body, and conclusion.
- Refine transition sentences to enhance flow between paragraphs and sections.
- Infuse your own insights, analysis, and voice to make it original and reflective of your understanding.
- Match the tone and style to your writing and the assignment's requirements.
- Harmonize word choice and sentence structure for clarity and readability, ensuring the text sounds cohesive and as if it were written by one author.
- AI-generated text is often generic and lacks specific arguments. Add specific examples, case studies, and evidence to support your arguments.
- Remove redundancies and expand the text with deeper analysis, critical thinking, and synthesis of ideas.
- Enrich the AI-generated content with additional research and properly cite any borrowed ideas, data, or phrases, including those from the AI tool.

Tools and ethics: Use additional tools like grammar checkers (for example, Grammarly, DeepL) and citation managers to refine your work. Ensure that your work is original, not just a rephrased AI-generated version. Begin by writing some text yourself and use AI tools to refine it in specific ways. Conduct additional research to find more sources, perspectives, and information.

Relying solely on AI tools for research is not advisable. However, by focusing on accuracy, clarity, personalization, depth, grammar, and ethics, you can turn AI-generated content into high-quality, original work.

2. Data Pitfalls

💡 Recommended reading

Please read “Quantitative Methods” by [Huber \[2024a\]](#). In particular, chapter “Identification” can help to understand how causes of effects can be identified empirically.

The following sections highlight common pitfalls when dealing with data. Being aware of these pitfalls helps build best practices to maintain the integrity of analyses and visualizations. The structure and content align with the excellent book by [Jones \[2020\]](#). I recommend reading this book.

2.1. Epistemic Errors

Learning objectives

- Understand the concept of epistemic errors and their impact on data analysis.
- Identify and describe common types of epistemic errors.
- Apply strategies to avoid epistemic errors in your data work.

Epistemic errors occur when there are mistakes in our understanding and conceptualization of data. These errors arise from cognitive biases, misunderstandings, and incorrect assumptions about the nature of data and the reality it represents. Recognizing and addressing these errors is crucial for accurate data analysis and effective decision-making.

One significant type of epistemic error is the data-reality gap, which refers to the difference between the data we collect and the reality it is supposed to represent. For instance, a survey on customer satisfaction that only includes responses from a self-selected group of highly engaged customers may not accurately reflect the overall customer base. To avoid this pitfall, it is essential to ensure that your data collection methods are representative and unbiased, and to validate your data against external benchmarks or additional data sources.

The Data-Reality Gap

The difference between the data we collect and the reality it is supposed to represent.

Another common epistemic error involves the influence of human biases during data collection and interpretation. Known as the all too human data error, this occurs when personal biases or inaccuracies affect the data. An example would be a researcher’s personal bias influencing the design of a study or the interpretation of its results. To mitigate this, implement rigorous protocols for data collection and analysis, and consider using double-blind studies and peer reviews to minimize bias.

2. Data Pitfalls

All Too Human Data

Errors introduced by human biases or inaccuracies during data collection and interpretation.

Inconsistent ratings can also lead to epistemic errors. This happens when there is variability in data collection methods, resulting in inconsistent or unreliable data. For example, different evaluators might rate the same product using different criteria or standards. To avoid this issue, standardize data collection processes and provide training for all individuals involved in data collection to ensure consistency.

Inconsistent Ratings

Variability in data collection methods that leads to inconsistent or unreliable data.

The black swan pitfall refers to the failure to account for rare, high-impact events that fall outside regular expectations. Financial models that did not predict the 2008 financial crisis due to the unexpected nature of the events that led to it are an example of this error. To prevent such pitfalls, consider a wide range of possible outcomes in your models and incorporate stress testing to understand the impact of rare events.

The Black Swan Pitfall

The failure to account for rare, high-impact events that fall outside the realm of regular expectations.

Falsifiability and the God pitfall involve the tendency to accept hypotheses that cannot be tested or disproven. This error might occur when assuming that a correlation between two variables implies causation without the ability to test alternative explanations. To avoid this, ensure that your hypotheses are testable and that you actively seek out potential falsifications. Use control groups and randomized experiments to validate causal relationships.

Falsifiability and the God Pitfall


The tendency to accept hypotheses that cannot be tested or disproven.

To avoid epistemic errors, critically assess your assumptions, methodologies, and interpretations. Engage in critical thinking by regularly questioning your assumptions and seeking alternative explanations for your findings. Employ methodological rigor by using standardized and validated methods for data collection and analysis. Engage with peers to review and critique your work, providing a fresh perspective and identifying potential biases. Finally, stay updated with the latest research and best practices in your field to avoid outdated or incorrect methodologies.

Understanding and addressing epistemic errors can significantly improve the reliability and accuracy of your data analyses, leading to better decision-making and more trustworthy insights.

Exercise 2.1.

Figure 2.1.: Bananas in various stages of ripeness



Source: Jones [2020, p. 33]

- a) Rate the ripeness level of the bananas pictured by Figure 2.1. Compare your assessment to that of a colleague and discuss any differences in your ratings. What might account for the variance in perception of the bananas' ripeness between you and your colleague?
- b) Specify how you rated the second and the last bananas on the ripeness scale?
- c) Upon reevaluation, it appears that the second and the last bananas are identical in ripeness. How would you justify your initial decision now? This scenario underscores an important lesson for interpreting polls and surveys: it illustrates how subjective assessments can lead to variance in results. It highlights the necessity of ensuring clarity and consistency in the criteria used for evaluations to minimize subjective discrepancies.

2.2. Technical trespasses

The second pitfall refers to the mishandling of data through a lack of technical expertise or the misapplication of statistical methods. It can lead to significant distortions in data analysis and result in flawed decisions. This might include using inappropriate statistical tests, excessive model complexity, or improper data manipulation. Understanding essential statistics and maintaining data integrity are vital to prevent these errors. Business leaders need to ensure that their teams have the requisite technical knowledge and to foster a culture of cooperation with data science experts to safeguard against these common technical blunders. This will ultimately support more reliable and actionable insights for strategic decision-making.

Here are some examples things that frequently causes wrong conclusions when analyzing data:

Misinterpreting correlation as causation

A classic mistake is interpreting a correlation between two variables as evidence that one causes the other. Without proper experimental design or further statistical analysis, such claims can lead to ill-informed business decisions.

2. Data Pitfalls

Inappropriate use of average

Averages can be misleading when data is skewed or has outliers. For example, using the mean salary in a company where the CEO's salary is drastically higher than the rest of the employees can give a false impression of the typical compensation.

Neglecting the test assumptions

Applying statistical tests without checking the necessary assumptions can invalidate the results. For instance, using a t-test on data that is not normally distributed or has unequal variances between groups might lead to incorrect conclusions.

Overfitting predictive models

Creating a model that performs extremely well on historical data but fails to predict future outcomes undermines its utility. Overfit models are often too complex, capturing noise rather than underlying trends in the training data.

Ignoring the effect of sample size

Small samples can produce volatile or non-representative results; business decisions made on such shaky ground may not stand up when scaled to the full population or market.

Cherry-picking data

Selecting data that supports a preconceived hypothesis, while ignoring data that doesn't, compromises the objectivity of the analysis and can mislead stakeholders about the reality of the situation.

Improper data visualization

Misrepresenting data through poor chart choices or scales can mislead viewers. For example, truncating the y-axis can exaggerate trends, and using pie charts for too many categories can confuse rather than clarify.

Exercise 2.2. Survivorship bias

Read "How Successful Leaders Think" by Roger [Martin \[2007\]](#) and the chapter "Identification" of "Quantitative Methods" by [Huber \[2024a\]](#).

Here is a summary of [Martin \[2007\]](#) taken from the [Harvard Business Review Store](#):

In search of lessons to apply in our own careers, we often try to emulate what effective leaders do. Roger Martin says this focus is misplaced, because moves that work in one context may make little sense in another. A more productive, though more difficult, approach is to look at how such leaders think. After extensive interviews with more than 50 of them, the author discovered that most are integrative thinkers—that is, they can hold in their heads two opposing ideas at once and then come up with a new idea that contains elements of each but is superior to both. Martin argues that this process of consideration and synthesis (rather than superior strategy or faultless execution) is the hallmark of exceptional businesses and the people who run them. To support

2. Data Pitfalls

his point, he examines how integrative thinkers approach the four stages of decision making to craft superior solutions. First, when determining which features of a problem are salient, they go beyond those that are obviously relevant. Second, they consider multidirectional and nonlinear relationships, not just linear ones. Third, they see the whole problem and how the parts fit together. Fourth, they creatively resolve the tensions between opposing ideas and generate new alternatives. According to the author, integrative thinking is an ability everyone can hone. He points to several examples of business leaders who have done so, such as Bob Young, co-founder and former CEO of Red Hat, the dominant distributor of Linux open-source software. Young recognized from the beginning that he didn't have to choose between the two prevailing software business models. Inspired by both, he forged an innovative third way, creating a service offering for corporate customers that placed Red Hat on a path to tremendous success.


a) Discuss the concepts introduced by [Martin \[2007\]](#) critically:

- Does he provide evidence for his ideas to work?
- Is there a proof that his suggestions can yield success?
- Is there some evidence about whether his ideas are superior to alternative causes of action?
- What can we learn from the article?
- Does his argumentation fulfill highest academic standards?
- What is his identification strategy with respect to the *causes of effects* and the *effects of causes*?
- [Martin \[2007, p. 81\]](#) speculates:

“At some point, integrative thinking will no longer be just a tacit skill (cultivated knowingly or not) in the heads of a select few.”

b) If teachers in business schools would have followed his ideas of integrative thinkers being more successful, almost 20 years later, this should be the dominant way to think as a business leader. Is that the case? And if so, can you still gain some competitive advantage by thinking that way?

Figure 2.2.: Distribution of bullet holes in returned aircraft



Source: [Martin Grandjean \(vector\)](#), [McGeddon \(picture\)](#), [Cameron Moll \(concept\)](#), [CC BY-SA 4.0, Link](#)

- c) Figure 2.2 visualizes the distribution of bullet holes in aircraft that returned from combat in World War II. Imagine you are an aircraft engineer. What does this picture teach you?
- d) Inform yourself about the concept of survivorship bias explained in [Wikipedia \[2024\]](#).
- e) In [Martin \[2007\]](#), the author provides an example of a successful company to support his management ideas. Discuss whether this article relates to survivorship bias.

2.3. Mathematical Miscues

This pitfall deals with the errors that can occur when mathematical reasoning is flawed or when calculations are incorrectly applied within data analysis. This encompasses a range of issues, from simple arithmetic mistakes to more complex misunderstandings of mathematical principles that underlie data models. It might involve misjudging how certain metrics are calculated, which could significantly alter an analysis's outcomes.

Additionally, this pitfall includes a lack of understanding the nuanced relationships within mathematical functions that can result in poor model performance or misestimation of forecasts and trends. Business leaders should be keenly aware that even advanced data analytics rely on the precise and proper application of mathematics. Therefore, ensuring precise calculations and fostering a thorough understanding of the mathematical techniques used in data analysis is vital to avoid undermining data-driven business decisions with slip-ups in mathematical reasoning.

Here is an incomplete list of things that are often misused:

Misuse of averages

Arithmetic Mean vs. Median: Incorrectly using the arithmetic mean without considering outliers or skewed distributions, which can misrepresent the central tendency of data.

Weighted Averages: Failing to appropriately weight averages based on sample sizes or importance of data points, leading to biased results.

Improper scaling and normalization

Scaling data improperly without considering the original distribution or range, which can distort comparisons and statistical analyses. Improper normalization techniques that fail to account for the underlying data structure or variance, leading to skewed interpretations.

Misapplication of probability and statistics

Misunderstanding statistical significance and confidence intervals, leading to incorrect conclusions about the reliability of findings. Inappropriate use of probability distributions or assumptions, which can invalidate statistical tests or predictions.

Errors in calculating percentages and ratios

Incorrectly calculating percentages or ratios without ensuring consistent bases or understanding the context, leading to misrepresentation of trends or relationships.

2.4. Statistical slipups

This pitfall is about common errors when applying statistical methods. These errors include misinterpreting statistical significance, sampling biases, ignoring multicollinearity in regression analysis, misapplying hypothesis tests, and failing to adjust for multiple comparisons.

For example, misinterpreting a small p-value as evidence of causation rather than correlation can lead to misleading conclusions. Similarly, using a convenience sample that does not accurately represent the population can skew results and compromise the generalizability of findings. In regression analysis, overlooking multicollinearity among predictor variables can result in inaccurate assessments of their individual impacts.

2.4.1. Descriptive debacles

This pitfall refers to errors in data analysis stemming from inadequate or misleading use of descriptive statistics. This bias occurs when analysts present summary statistics that are incomplete, misleading, or not properly contextualized. It also includes misrepresenting data through inappropriate visualizations or failing to provide sufficient detail in describing data characteristics.

In practical terms, this bias leads to misunderstandings or incorrect interpretations of data. For example, using measures like mean or median without considering data variability can oversimplify complex datasets. Similarly, misleading graphs or charts may obscure important trends or relationships within the data.

To mitigate this bias, analysts should employ a range of appropriate descriptive statistics tailored to data types and research questions. They should ensure visual representations accurately depict data characteristics without distorting interpretations. Providing adequate context and explanations helps stakeholders understand data limitations and implications more effectively.

2.4.2. Inferential infernos

Frequently, researchers draw erroneous or exaggerated conclusions from statistical analyses. This bias manifests when statistical significance is misinterpreted, leading to unwarranted claims of causation or generalization beyond the scope of the data.

Researchers may overlook alternative explanations or fail to consider confounding variables, thereby overstating the implications of their findings. To mitigate this bias, it is essential to apply statistical tests accurately, interpret results cautiously, and acknowledge the limitations and uncertainties of the data.

Recommended reading

Read [Wysocki et al. \[2022\]](#) which is freely available [here](#). [Here](#) you find a good summary of the paper.

2.4.3. Slippery sampling

Systematic errors can be introduced in research studies due to flawed or biased sampling methods. This bias occurs when the sample selected for study does not accurately represent the target population, leading to skewed or misleading conclusions.

In practice, researchers may inadvertently use convenience sampling methods or fail to account for selection biases, resulting in a sample that is not representative of the broader population. As a result, findings based on such samples may not generalize well beyond the specific sample group, undermining the external validity of the study.

Recommended reading

Read chapter [Sampling](#) of [Huber \[2024a\]](#).

2.4.4. Insensitivity to sample size

This bias refers to the tendency to overlook or underestimate the impact of sample size on the reliability and validity of study results. This bias occurs when analysts or researchers fail to recognize that smaller sample sizes may produce more variable or less representative data, leading to less reliable conclusions.

In practical terms, studies with smaller sample sizes are more prone to random fluctuations and may not accurately reflect the characteristics of the larger population. Consequently, findings based on small samples may lack statistical power and may not be generalizable to broader contexts or populations.

Here is an incomplete summary of common statistical slipups:

Misunderstanding statistical significance

Incorrect interpretation of p-values and confidence intervals can lead to erroneous conclusions about the strength of relationships in data.

Sampling errors

Biased sampling or inadequate sample size selection can distort results and impair the generalizability of findings.

Ignoring multicollinearity

Failure to account for correlations between predictor variables in regression analysis can lead to unreliable coefficient estimates and misjudgment of variable impacts.

Misapplication of hypothesis tests

Improper formulation of hypotheses or disregarding test assumptions can invalidate statistical results and mislead interpretations.


Failure to adjust for multiple comparisons

Conducting multiple tests without adjusting significance levels increases the likelihood of false positives and misleading conclusions.

2.5. Analytical Aberrations

2.5.1. The intuition/analysis false dichotomy

This error stems from misconception that intuition and analytical thinking involves when trying to use it as approaches to decision-making. In reality, effective decision-making often integrates both intuitive insights and analytical rigor. Intuition can provide valuable initial impressions or creative ideas, while analysis offers systematic evaluation and validation of hypotheses. Acknowledging and balancing these two aspects enhances decision-making processes by leveraging diverse perspectives and ensuring comprehensive consideration of factors influencing outcomes.

 Recommended reading

Read the chapter [Decision making basics](#) of [Huber \[2024c\]](#)

While intuition comes with multiple cognitive biases. Intuition remains valuable in decision-making for several reasons. Firstly, it often draws on subconscious processing of vast amounts of information, leading to quick insights that can guide effective decisions. Secondly, intuition can fill gaps left by incomplete data, offering solutions when data is scarce or ambiguous. Thirdly, it allows for creativity and innovation by providing unconventional perspectives that data alone may not reveal. Fourthly, intuition serves as a valuable complement to analytical thinking, providing a holistic approach to problem-solving. Lastly, it facilitates swift decision-making in fast-paced environments where immediate action is crucial. Integrating intuition with data-driven approaches enhances decision-making by leveraging both analytical rigor and intuitive wisdom.

2.5.2. Exuberant extrapolations

Exuberant Extrapolations refer to the tendency to make overly optimistic or exaggerated predictions based on limited data or trends. This pitfall occurs when analysts or decision-makers extend current trends far into the future without considering potential deviations or external factors that could influence outcomes. It can lead to unrealistic expectations, flawed projections, and poor decision-making due to overconfidence in extrapolated outcomes. Mitigating this pitfall requires cautious interpretation of trends, consideration of uncertainties, and validation through robust analysis and scenario planning to ensure more realistic and reliable forecasts.

Exercise 2.3. Past and present

Again, let us consider “How Successful Leaders Think” by Roger [Martin \[2007\]](#).

Discuss the insights that he puts forward. Can we really learn from the experience of successful leaders?

2.5.3. Ill-advised interpolations

Ill-advised interpolations occur when analysts or decision-makers infer trends or make assumptions between data points without sufficient evidence or justification. This pitfall involves filling gaps in data with speculative assumptions or linear extrapolations that do not accurately reflect the underlying patterns or relationships. It can lead to misleading interpretations, erroneous conclusions, and misguided decision-making based on incomplete or flawed information.

2.5.4. Funky forecasts

This pitfall arises when forecasters fail to account for changing circumstances, unexpected events, or complex interactions among factors influencing future trends. It can lead to unreliable projections, erroneous planning, and ineffective decision-making based on overly simplistic or deterministic forecasts. To mitigate this pitfall, analysts should use robust forecasting techniques, incorporate scenario analysis to account for uncertainties, and continuously update forecasts based on new information and changing conditions.

2.5.5. Moronic measures

Moronic measures refer to the use of inappropriate or inadequate metrics for evaluating performance or assessing outcomes. This pitfall occurs when decision-makers rely on metrics that do not accurately capture the intended goals or fail to account for relevant factors influencing performance. It can lead to misleading assessments, misguided resource allocation, and ineffective strategies due to a lack of alignment between chosen metrics and organizational objectives. To mitigate this pitfall, organizations should carefully select metrics that are relevant, measurable, and aligned with strategic priorities. They should also consider using complementary metrics and qualitative assessments to provide a comprehensive understanding of performance outcomes. By using appropriate measures and periodically reviewing their relevance, organizations can enhance decision-making and ensure that performance evaluations contribute meaningfully to achieving desired outcomes.

2.6. Graphical Gaffes

Recommended reading

Read the chapter [Visualize data](#) of [Huber \[2024b\]](#) and the literature that you find cited therein.

2.6.1. Challenging charts

Pitfalls in data visualization are pretty common. Charts are frequently poorly designed, misleading, or difficult to interpret. Actually, it is simply difficult to effectively communicate insights with graphical visualizations of data due to cluttered layouts, inappropriate chart types, distorted scales, or ambiguous labeling. It can lead to confusion, misinterpretation of data trends, and inaccurate decision-making based on flawed visual presentations. To mitigate this pitfall,

analysts should adhere to best practices in data visualization, such as using clear and simple designs, choosing appropriate chart types for the data being presented, ensuring accurate scaling and labeling, and providing contextual information to aid interpretation.

2.6.2. Data dogmatism

Data Dogmatism refers to the pitfall of rigidly adhering to data-driven decisions without considering broader context, qualitative insights, or expert judgment. This pitfall occurs when organizations or decision-makers prioritize data alone as the sole basis for decision-making, ignoring nuanced factors that may influence outcomes. It can lead to overly deterministic decisions, missed opportunities, and ineffective strategies that fail to account for human expertise, market dynamics, or unforeseen variables. To mitigate this pitfall, organizations should adopt a balanced approach that integrates data-driven insights with qualitative analysis, expert opinions, and contextual understanding. By recognizing the limitations of data and embracing a more holistic decision-making framework, organizations can enhance flexibility, innovation, and strategic agility in response to complex challenges.

2.6.3. The optimize/satisfice false dichotomy

This dichotomy refers to the misconception that decision-making strategies must exclusively prioritize either optimizing for the best possible outcome or satisficing by accepting a satisfactory but not necessarily optimal solution. This pitfall arises when decision-makers perceive a dichotomy between these approaches, potentially overlooking opportunities for adaptive strategies that blend elements of both optimization and satisficing. It can lead to suboptimal decisions, missed opportunities for innovation, or inefficient resource allocation due to rigid adherence to a singular decision-making approach. To mitigate this pitfall, organizations should adopt a flexible and adaptive decision-making framework that considers both optimizing for excellence where feasible and satisficing when constraints or uncertainties warrant. By embracing a nuanced approach that balances ambition with pragmatism, organizations can enhance resilience, innovation, and effectiveness in achieving strategic objectives.

2.7. Design Dangers

Recommended reading

Read the chapter [Visualize data](#) of [Huber \[2024b\]](#) and the literature that you find cited therein.

This section encompasses pitfalls related to the design and implementation of data-driven projects or experiments. It includes errors such as biased sampling methods, inadequate experimental design, or flawed data collection protocols.

2.7.1. Confusing Colors

In data visualization the choice of colors in charts matter. When colors are poorly contrasted, misaligned with data meaning, or inconsistently applied across visual elements, it may be difficult to distinguish different data categories or trends. The wrong colors can lead to misunderstandings and inaccurate conclusions. And by the way, many people are color blind.

2.7.2. Usability uh-ohs

This pitfall occurs when systems lack intuitive interfaces, are overly complex, or fail to meet user needs and expectations. It can result in frustration, inefficiencies, or resistance to adoption among users, leading to underutilization of valuable data resources. To mitigate this pitfall, organizations should prioritize user-centric design principles, conduct usability testing with diverse user groups, and solicit feedback to refine interfaces and functionalities. By enhancing usability, organizations can improve user satisfaction, facilitate smoother workflows, and maximize the utility of data-driven tools for informed decision-making.

2.7.2.1. Good visualizations are discoverable and understandable

Good visualizations are designed to be easily discovered and understood by users. They prioritize clarity and accessibility, allowing users to quickly grasp the presented information without ambiguity or unnecessary complexity.

2.7.2.2. Don't blame people for getting confused or making errors

Users should not be blamed for confusion or errors in understanding visualizations. Instead, designers should strive to create intuitive and user-friendly designs that minimize cognitive load and facilitate accurate interpretation of data.

2.7.2.3. Designing for pleasure and emotion is important

Effective visualizations consider the emotional and aesthetic impact on users. Design elements that evoke positive emotions and engage users can enhance comprehension and retention of information, making the visualization more impactful and memorable.

2.7.2.4. Complexity is good, confusion is bad

Complexity in visualizations can enrich understanding by presenting detailed information and relationships. However, complexity should not lead to confusion or overwhelm users. Designers must balance complexity with clarity to ensure that insights are communicated effectively without causing cognitive overload.

2.7.2.5. Absolute precision isn't always necessary

Visualizations do not always need to convey absolute precision in every detail. Depending on the context and audience, approximations or visual representations that convey trends and patterns effectively may be sufficient. Designers should prioritize meaningful insights over exhaustive accuracy in all data points.

3. How to use R for data science

The programming language R is one of the major tools to do data science. I wrote some lecture notes on [How to use R for data science](#) [Huber, 2024b].


Please read these notes.

Tip

We have discussed the following chapters 1, 2, 3, 5, 6, 7, 8 and some exercises of chapter 9. In particular, we have put some emphasise on chapter 6, sections 7.2-7.4, and chapter 8. Moreover, we have explicitly studied the exercises *Consumer prices over time*, *Load the Stata dataset “auto” using R*, *Convergence*,

4. Collaborating with Git and GitHub

Figure 4.1.: The FINAL.doc problem



Source: phdcomics.com

4.1. Introduction

Git is open-source software for version control. It allows developers to track and manage changes to their codebase and files. Users can access a comprehensive history of their project and revert to previous versions of their data if necessary. It helps to overcome the *FINAL.doc problem* depicted in Figure 4.1.



Source: <https://github.com/about> as of April 2024

GitHub is an incredibly popular (see statistics in Figure 4.2) online platform that implements Git's capabilities by providing a web interface for collaboration.

While you can use Git and GitHub independently, most developers integrate it with GitHub for enhanced project management and collaboration. This combination helps maintain local and remote copies of a project, facilitating teamwork and data backup as GitHub is sort of a

backup as data loss at your local machine do not matter if you have a remote version saved on GitHub.

Git and GitHub support simultaneous multi-user access, unlike systems that are optimized for single-user like Dropbox.

4.2. Install Git

To install the version control system Git, follow the instructions [here](#).

Figure 4.3.: Memorizing six git commands

The popular approach to version control



ORLY?

@ThePracticalDev

Source: [DEV Community on GitHub](#)

Familiarize yourself with Git by using the resources available [here](#). Specifically, work through the resources listed in the box below. Although Git may appear complex, it is generally not too challenging for most users. Many people use Git primarily to track their work and to host and share files conveniently with just a handful of commands. While Git is a robust system with many capabilities, you don't need to memorize all the commands (see Figure 4.3). In fact, you typically use only a few basic ones as shown in Table 4.1.

In the upcoming sections, I will demonstrate some use cases both in the terminal Section 4.3 and within RStudio Section 4.4. In Section 4.5, I show how to contribute to a repository using Git and GitHub.

💡 Learning resources

Plenty books and tutorial exist that introduce Git and GitHub. I'd like to highlight the following sources:

- The book comprehensive book [Happy Git and GitHub for the useR](#) by [Bryan](#)
- The much shorter book [[Version Control with Git and GitHub](#)] by [Halbritter and Telford](#)
- The online tutorial [How to Use Git/GitHub with R](#) of David Keyes who explains

in short videos how to setup Git and GitHub in RStudio using among others the `usethis` package.

Table 4.1.: Most important git commands

Git Command	Description
<code>git init</code>	Initialize a new Git repository in the current directory.
<code>git clone <url></code>	Clone a repository from a remote URL to your local machine.
<code>git add <file></code>	Add a specific file to the staging area in preparation for committing.
<code>git add .</code>	Add all changed files in the current directory to the staging area.
<code>git commit -m "message"</code>	Commit the staged changes to the repository with a descriptive message.
<code>git status</code>	Display the status of the working directory and staging area.
<code>git push <remote> <branch></code>	Push committed changes in your local branch to the remote repository.
<code>git pull <remote> <branch></code>	Pull changes from the remote repository into your current branch and merge them.
<code>git branch <name></code>	Create a new branch with the specified name.
<code>git checkout <branch></code>	Switch to another branch and update the working directory.
<code>git merge <branch></code>	Merge a specified branch into the current branch.

4.3. Using Git from the terminal

Figure 4.4.: Three git commands you really need

```
git add .
git commit -m "message"
git push
```

This tutorial will guide you through the basic Git operations using the Bash command line, commonly referred to as the terminal. Essentially, it focuses on the three Git commands illustrated in Figure 4.4.

4.3.1. Configuring Git

Before you start using Git, you need to configure your Git environment. Set your username and email address with these commands:


```
git config --global user.name "Your Name"
git config --global user.email "your.email@example.com"
```

4.3.2. Initializing a Repository

To create a new Git repository, use the `git init` command in the directory you want to version control:

```
cd /path/to/a/directory
mkdir my_project
cd my_project
git init
```

In case you are not familiar with using the terminal please consider Table 4.2 where I introduce the most basic commands that we use. For example, with `cd` you can change your directory and with `mkdir` you create a new directory. If you are not familiar with the file system of your computer please read the section *Navigating the file system* of Huber [2024b]. With `git init` you initialize the directory to be a git repository. This will create a hidden folder “.git” in which Git keeps track of all your changes.

i Most common bash commands

Table 4.2.: Most common bash commands

Bash Command (macOS/Linux)	Windows Command Prompt Equivalent	Description
<code>pwd</code>	<code>cd</code>	Prints the current directory's path.
<code>ls</code>	<code>dir</code>	Lists all files and directories in the current directory.
<code>cd</code>	<code>cd</code>	Changes the directory.
<code>mkdir</code>	<code>mkdir</code>	Creates a new directory.
<code>rmdir</code>	<code>rmdir</code>	Removes an empty directory.
<code>touch</code>	<code>copy nul</code>	Creates a new empty file or updates an existing file's timestamp.
<code>rm</code>	<code>del</code> or <code>erase</code>	Removes files. <code>rmdir /s</code> is used for directories.
<code>cp</code>	<code>copy</code>	Copies files or directories.
<code>mv</code>	<code>move</code>	Moves or renames files or directories.
<code>echo</code>	<code>echo</code>	Displays a line of text/string.
<code>cat</code>	<code>type</code>	Concatenates and displays the content of files.
<code>grep</code>	<code>find</code> or <code>findstr</code>	Searches for patterns in files.

4.3.3. Staging Changes

To track changes in your repository, you need to stage them using the `git add` command. To stage a single file:

```
git add filename.txt
```

To stage all changes in the directory:

```
git add .
```

4.3.4. Committing Changes

After staging, you can commit it to the repository. A commit records changes to the repository and must include a message describing what changed:

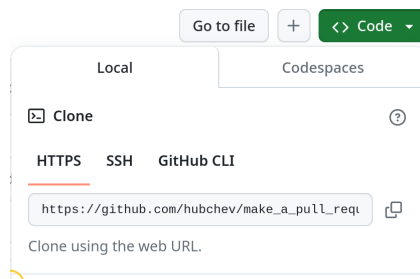
```
git commit -m "A message"
```

4.3.5. Pushing Changes

To share your commits with others or store them in a remote repository (GitHub), use `git push`. A prerequisite here is that you need to be connected to a remote repo. Therefore, you must add a remote repository by copying the URL of the GitHub repo as shown in Figure 4.5. Then you can add the remote repository and push it to the repo with these lines of code:

```
git remote add origin https://github.com/username/repository.git
git push -u origin main
```

Figure 4.5.: Copy the https URL of your repo



4.3.6. Undo changes

With `git reset` and `git revert` you can go back in time and undo specific changes, respectively. For example, with

```
git log
git reset --hard <commit_id_hash>
```

you can view the commit history and find the hash identifier of the commit to move the HEAD pointer to that commit. This effectively removes all commits after commit you choose from the current branch's history. Be cautious when using `git reset -hard` as it discards all changes made after the specified commit. Make sure you have backups or are certain you want to discard these changes before proceeding.

With

```
git revert <commit_id_hash>
```

you revert the changes introduced by that commit only. It will create a new commit that undoes the changes made in commit chosen while keeping the other commits that may have followed the chosen commit intact. It's a safer approach compared to `git reset --hard`, as it preserves the commit history and allows you to selectively undo changes without affecting the rest of the commits.

4.4. Using Git from RStudio

Integrating Git with RStudio enhances your project management by utilizing version control directly within the IDE. Here's how you can set up and use Git in RStudio using R code.

4.4.1. Set up Git in RStudio

First, ensure the `usethis` package is installed and loaded:

```
if (!require(pacman)) install.packages("pacman")
pacman::p_load(usethis)
```

Configure your Git settings in RStudio:

```
use_git_config(user.name = "Your Name", user.email = "Your@email.com")
```

You can change the configuration of your user name and email using the `edit_git_config()` function.

Start a new project in RStudio, which will also initialize a Git repository:

```
create_project("~/Music/")
use_git()
```

After restarting RStudio, you will notice a Git tab in the top right panel, indicating that Git is now active for your project.

4.4.2. Connecting RStudio Projects with GitHub repositories

To connect your RStudio project with GitHub, you need a Personal Access Token (PAT) on GitHub. If you haven't one already, you can use the `create_github_token()` function from `usethis` package, and store the PAT securely with `gitcreds_set` from the `gitcreds` package:

```
if (!require(pacman)) install.packages("pacman")
pacman::p_load(usethis gitcreds)
create_github_token()
gitcreds::gitcreds_set()
```

Now, the procedure depends on whether the project has been initialized on your local machine and you want to create a repo on GitHub, or the repo already exists on GitHub and you want to connect that remote repo with your local PC. Both ways are described below.

4.4.2.1. Project exists on RStudio first

After initializing Git in your project, use the `use_github()` function from `usethis` to create a new GitHub repository and connect it directly:

```
use_github()
```

This creates a repo on your GitHub account.

4.4.2.2. Project exists on GitHub first

Alternatively, suppose you have created a repository on GitHub first, then start a new project in RStudio using the version control option, specifying your new repository's URL. Just click **File > New Project > Version Control** and then link the GitHub repo by putting the URL into the respective box of the menu. See Figure 4.5 how to get the URL of a repo.

4.5. Make a contribution using Git and GitHub

This is a guide for beginners on how to make a contribution using Git and GitHub. If you are looking to make your first contribution, follow the steps below.

Watch

this [video](#) where I do all the following steps in real time. It takes about 15 minutes.

1. Create an account on GitHub

It is for free and should just take some minutes.

2. Install Git

[Here](#) is a tutorial on how to set up Git.

3. Fork [this](#) repository

Click on the fork button (see Figure 4.6) on the top of this page: https://github.com/hubchev/make_a_pull_request. This will create a copy of this repository in your account.



4. Clone the forked repository

Go to your GitHub account, open the forked repository, click on the code button and then click the *copy to clipboard* icon, see Figure 4.5.

Then, open a terminal and run the following git command:

```
git clone "url you just copied"
```

4. Collaborating with Git and GitHub

where “url you just copied” (without the quotation marks) is the url to this repository (your fork of this project). See the previous steps to obtain the url.

For example:

```
git clone https://github.com/this-is-you/make_a_pull_request.git
```

where `this-is-you` is your GitHub username. Here you’re copying the contents of the `first-contributions` repository on GitHub to your computer.

5. Create a branch

Change to the repository directory on your computer (if you are not already there):

```
cd make_a_pull_request
```

Now create a branch using the `git switch` command:

```
git switch -c your-new-branch-name
```

For example:

```
git switch -c add-Stephan-Huber
```

6. Make changes.

Now open the `I_am_a_data_scientist.md` file in a text editor. (You find this file in the repository.) Add your name, your GitHub account and the project you are working on. You can put it anywhere in between. Now, save the file.

If you go to the project directory and execute the command `git status`, you’ll see there are changes.

7. Add changes (staging). Add those changes to the branch you just created using the `git add` command:

```
git add .
```

8. Commit changes. Now commit those changes using the `git commit` command:

```
git commit -m "Add your-name to the list"
```

replacing `your-name` with your name.

9. Use Git Bash. Open Git Bash and set your email and your nickname on GitHub:

```
git config --global user.name "FIRST_NAME LAST_NAME"  
git config --global user.email "MY_NAME@example.com"
```

10. Push changes to GitHub.

Push your changes using the command `git push`:

4. Collaborating with Git and GitHub

```
git push -u origin your-new-branch-name
```

replacing `your-new-branch-name` with the name of the branch you created earlier.

If you get any errors while pushing that refers to authentication failed something, go to [GitHub's tutorial](#) on generating and configuring an SSH key to your account. Alternatively, you can watch this [YouTube tutorial](#)

11. Submit your changes for review on GitHub.

If you go to your repository on GitHub, you'll see a `Compare & pull request` button. Click on that button.

Now submit the pull request.

Soon I'll be merging all your changes into the main branch of this project. You will get a notification email once the changes have been merged.

Congrats! You just completed the standard *fork -> clone -> edit -> pull request* workflow that you'll often encounter as a contributor!

5. Markdown, Quarto, and R Markdown

Verbal and non-verbal communication are crucial in business. This section focuses on writing and publishing texts, excluding aspects like body language and writing skills. I will introduce some tools commonly used by data scientists for writing and publishing their work, such as Markdown, RMarkdown, and Quarto. Unlike applications like Microsoft Word or Apple Pages, these tools use code to generate text. This concept may be unfamiliar to those who grew up after Windows 95; I will provide justification for its use in the following sections.

One notable advantage of a code-based approach to writing text is its seamless integration with version control systems like Git and platforms like GitHub. These tools are essential for most data science collaborations. Mastering them can greatly enhance your efficiency and make your presentations more impactful, even if you are not directly involved in data science.

What is Quarto?

Quarto, a modern documentation system, is an excellent choice for writing, especially for projects that require rigorous data analysis, visualization, and reproducibility. This tutorial will guide you through producing various forms of text with Quarto. You can write reports, articles, theses, books, websites and many more with Quarto.

i Quarto and R Markdown

Quarto is a relatively new tool and can be considered as a successor to R Markdown. Most R Markdown documents are compatible with Quarto. However, Quarto offers some improved functionality over R Markdown, which enhances user-friendliness. A detailed overview of the differences and similarities between the two can be found in [this article](#). For an introduction to R Markdown see Section 5.4.

5.1. Why Markdown and Quarto?

5.1.1. No-code vs. code-based writing applications

Students often use Microsoft Word, Apple Pages, or LibreOffice to write scientific texts. These word processing programs operate on the “What You See Is What You Get” (WYSIWYG) principle, displaying the document layout as you type. While this principle and its corresponding applications are widespread and may seem indispensable to many, this is far from true. Alternatives such as LaTeX, Markdown, R Markdown, and Quarto offer significant advantages. Many professional scientists and publishers prefer these alternatives for good reason. A large number of doctoral theses and scientific papers are authored using LaTeX, and nearly all publishers and editors work with code-based solutions that do not follow the WYSIWYG principle.

With code-based alternatives, layout specifications are either placed at the beginning of the text or embedded within the main text itself. The final document is only visible after converting (also called compiling or rendering) it into a format such as PDF. This may initially seem unusual

and less intuitive than a WYSIWYG interface, but the most intuitive solution is not necessarily the best or simplest. My experience supervising numerous student papers has shown that the intuitive features of MS Word and Pages often become time-consuming over the medium to long term and fail to adequately support users in avoiding errors when writing scientific work. Students who choose code-based applications tend to experience less frustration and greater success—at least, this has been true for the papers I have supervised.

Code-based applications allow writers to focus on the actual writing process, as formatting and adherence to citation rules are largely automated by the software. The necessary initial investment in learning a tool like Quarto quickly pays off, resulting in noticeable improvements in the quality of scientific texts.

In the following subsections, I will first outline typical usage of WYSIWYG applications, then discuss the advantages of code-based text creation using Quarto as an example, and finally explain how to successfully write texts using Quarto.

5.1.2. Typical (mis)usage of WYSIWYG applications

The use of traditional word processing software like Microsoft Word or Apple Pages for writing academic papers is pervasive among students. While these programs are user-friendly for everyday writing projects, they create a significant additional workload to meet the demands of academic work.

One of the first problems is integrating literature. Correct formatting according to various citation guidelines is often counter-intuitive, and errors occur easily. This is particularly true if the citation and bibliography functions provided by the software are not used or are used incorrectly. Instead of utilizing external citation managers and investing time to learn how to use them, many students manually create citations and bibliographies, which typically leads to numerous small and sometimes larger errors that could be avoided.

Another weak point in student work is adherence to specific formatting requirements. Academic institutions and journals often require strict adherence to formatting guidelines, including the design of title pages, headers and footers, page margins, and heading hierarchies. Although Word and Pages offer templates and styles, they must be individually adapted for each document and often modified due to minor text changes. Making a format adjustment can become a major effort.

The inclusion of empirical results such as statistical data and graphics presents an additional hurdle. With Word and Pages, the process is frequently manual: research data must be exported from statistical software, saved as images, and then embedded in the document. If the data changes, this time-consuming process must be repeated, significantly increasing the workload and risk of errors.

5.1.3. Advantages of Quarto for writing text

Writing academic texts using traditional tools such as MS Word or Pages can be time-consuming and error-prone for students. In the following section, I introduce Quarto (or R Markdown), a modern alternative that offers several advantages:

- **Versatile output formats:** Quarto makes it effortless to generate different output formats. The same text can be rendered as a website (HTML), manuscript (PDF, DOCX), book (EPUB, PDF), or slides (PDF). This flexibility allows you to focus more on the content than the format.

- **Simplified formatting changes:** Specific templates can be used in Quarto, simplifying the process of making formatting changes.
- **Seamless literature integration:** Quarto handles citation rules compliance and integrates seamlessly with citation management systems, enabling researchers to manage literature references and bibliographies more efficiently and consistently than in Word.
- **Easy cross-referencing:** Creating cross-references to sections, tables, and figures is straightforward.
- **Direct data analysis and output generation:** Data analysis and output generation occur directly within Quarto, ensuring that displayed graphics and tables are always up-to-date. This eliminates the need for manual post-processing and guarantees the reproducibility of results.
- **Embedded data visualizations:** Researchers can embed data visualizations directly in the text without manual intermediate steps.
- **Efficient collaboration with version control:** Version control systems like Git make collaboration on academic documents more manageable. Changes can be tracked and integrated without relying on complex and conflict-prone comparison tools.

 Reading recommendation

For those interested, I recommend the online course [Introduction to Reproducible Publications with RStudio](#), which explains explicitly how to work in an empirically reproducible manner. A somewhat more compact introduction is offered by [Bauer and Landesvatter \[2023\]](#), and the authoritative work on the subject is by [Gandrud \[2020\]](#).

5.2. Markdown

Markdown is a lightweight markup language with plain-text formatting syntax. It is very popular because it is easy to learn. It's an essential skill for using Quarto effectively. Start by learning enough Markdown to structure your thesis, including headings, lists, links, and code blocks.

You can learn Markdown (not R Markdown!) in 10 minutes. Just go to www.markdowntutorial.com and work through the interactive lessons. I also recommend the introduction offered in the section [Markdown Basics](#) on quarto.org.

5.3. Quarto

 Recommended literature

Read [Telford \[2023\]](#): [Enough Markdown to Write a Thesis](#). This resource covers the basics and some advanced Markdown features that are useful for academic writing. More extensive resources on how to do things with Quarto can be found at quarto.org.

5.3.1. Introduction

To set up Quarto on your machine do the following:

- Install R and R Studio.

- Install Quarto as follows:

```
install.packages("quarto")
```

- Install the tinytex package to generate PDF files:

```
install.packages("tinytex")
tinytex::install_tinytex()
```

- It is also advisable to install additional packages that might be needed later:

```
if (!require(pacman)) install.packages("pacman")
pacman::p_load(knitr, rmarkdown, papaja)
```

Exercise 5.1. First Quarto document

- Open RStudio.
- Select “File” -> “New File” -> “Quarto Document” and then “Create.”
- Save the new file in an empty folder and set this folder as your working directory.
- Click “Render.”
- Visit the Markdown Basics website, add some Markdown to your document, and click “Render” again.
- Click the arrow next to the “Render” button. Here, you can select and generate other file formats. Give it a try.
- Consult the PDF Basics website and supplement your header with the information found there.
- Try citing the paper by [Huber and Rust \[2016\]](#), which you can find here, in your document.
 - Click on “Visual,”
 - Go to the place in the text where you want to cite the paper and select “Insert” -> “Citation.”
 - Search for the paper in the context menu using the corresponding DOI (<https://doi.org/10.1177/1536867X1601600209>) and insert it.
- To quote using APA Version 7 style, write the following in the YAML header:

```
csl: "https://www.zotero.org/styles/apa"
```

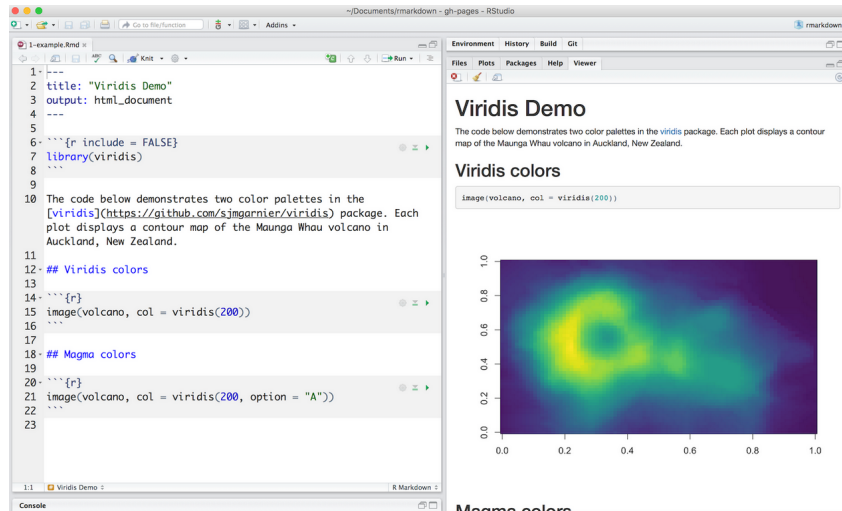
- Select a different citation style from www.zotero.org/styles. Then render the document again and observe the differences.

5.3.2. Create an APA compliant manuscript using Quarto

To create an APA compliant manuscript, it is recommended to use the *Quarto Extension* `apaquarto`. The process is described in detail [here](#). Using the template ensures that all APA rules are automatically considered. As APA allows a lot of leeway and every reviewer has specific preferences, `apaquarto` allows for manipulation of a variety of settings. For example, [the language can be changed](#) and the [general style of the document can be modified](#) in the Preamble (YAML header).

5.4. R Markdown

Figure 5.1.: Example of an R Markdown file




R Markdown provides an authoring framework for data science. You can use a single R Markdown file to transcript your work, run code, and generate high quality reports, books, websites, articles, theses, blogs, and many more (see Figure 5.1).

In contrast to Quarto (see Chapter 5), which is the more recent format, R Markdown is around for some time and hence there are uncountable resources to learn it. For example:

- The [R Markdown Cheatsheet](#) (see Figure 5.2) from Posit offers an overview on the most important features of R Markdown.

Figure 5.2.: R Markdown Cheatsheet from Posit



- The book *R Markdown Cookbook* by [Xie et al. \[2020\]](#) (see Figure 5.3) offers an introduction. The [online version of the book](#) is regularly updated and free of costs.
- The book *R Markdown: The Definitive Guide* by [Xie et al. \[2018\]](#) offers a comprehensive introduction. The [online version of the book](#) is regularly updated and free of costs.

Please watch the video [What is R Markdown?](#) and then study the [R Markdown tutorial from RStudio](#).

💡 Working directory in R Markdown

The working directory is by default set to the directory that contains the Rmd document. In case you want to use another directory you can do so by changing the working directory with `setwd()`. However, that is not persistent in R Markdown and only works for the

5. Markdown, Quarto, and R Markdown

Figure 5.3.: Xie et al. [2020]: R Markdown Cookbook





Figure 5.4.: Xie et al. [2018]: R Markdown: The Definitive Guide



current code chunk. After the code chunk has been evaluated, the working directory will be restored to the directory where the Rmd file is placed.

Exercise 5.2. Start Markdown and R Markdown

- a) You can learn Markdown (not R Markdown!) in 10 minutes. Just go to <https://www.markdwntutorial.com> and work through the interactive lessons.
- b) Now create your first R Markdown file in 3 minutes by doing the following:
 - click in RStudio on *File > New File > R Markdown*
 - click *OK*
 - look for a button entitled *Knit* and click it
 - save your file (it will be saved with *.Rmd* file extension)
- c) Play around with the file. For example, change the output format can you create a word file or a presentation. Play around with the code chunks. Add a picture that you find somewhere online.
- d) Set your working directory to the folder where you have saved your first Rmd-file. Can you come up with a way to generate different output format with just one function.

Exercise 5.3. R Markdown cite literature

- a) Create a new R Markdown file (*File > New File > R Markdown*), save the file in an empty folder, and knit it.
- b)
 - Make a new script with *File > New File > R Script*.
 - Go to <https://scholar.google.de/> and search for *osrmtime*.
 - Click on “*cite*” and “*BibTeX*”. Copy and paste everything that you see into your script and save the script as *lit.bib*. R Studio will ask you if you confirm the file type change. Click yes. Your *lit.bib* file should look like this:

```
@article{huber2016calculate,
  title={Calculate travel time and distance with OpenStreetMap
  data using the Open Source Routing Machine (OSRM)},
  author={Huber, Stephan and Rust, Christoph},
  journal={The Stata Journal},
  volume={16},
  number={2},
  pages={416--423},
  year={2016},
  publisher={SAGE Publications Sage CA: Los Angeles, CA}
}
```

- Add the text “*bibliography: references.bib*” to your YAML header of your R Markdown file so that it looks somehow like that:

```

---
title: "Untitled"
author: "Stephan Huber"
date: "`r Sys.Date()`"
output: html_document
bibliography: lit.bib
---

```

- Now you can cite the OSRMTIME paper with `@huber2016calculate` somewhere in the text of your R Markdown file.
- Knit the R Markdown file and you should see the paper cited and a reference list at the end of the html report.
- You can manipulate the citation style you can specify a CSL (Citation Style Language) file in the YAML header. For example the APA style can be chosen with:

```
csl: "https://www.zotero.org/styles/apa.csl"
```

Many more citation styles can be found on github.com/citation-style-language and on the [Zotero Style Repository](https://www.zotero.org/styles/).

Exercise 5.4. Preparing APA journal articles (`papaja`)

There is an easy way to write a manuscript that follows all the APA rules using the package `papaja` written by two psychologists from Cologne. Please read their [manual](#) and consider their [repository on GitHub](#).

Now, install and load the package:

```
install.packages("papaja")
library("papaja")
```

Then, click “File > New File > R Markdown” and choose the “APA-style manuscript” from the section “from template”. Knit the R markdown template and you will have a template for a APA manuscript.

Apart from the obvious adjustments, I recommend to make at least two general adjustments: Change `classoption` to “doc” and `linenumbers` to “no”.

Exercise 5.5. R Markdown template

Please follow the instructions below to access the file “23-09_ds-project-desc.Rmd” from my GitHub account:

1. Download the file from my GitHub account by clicking on the link provided here.
2. Save the file in your working directory.
3. Use the `knit` function to run the file, but be aware that it may not work properly at first. If you encounter any issues, troubleshooting may be required. Don’t worry, error messages will usually provide guidance to help you resolve the issue. Please note that the YAML header is sensitive to spacing, so be careful when setting it up to avoid breaking the code.
4. In the project template, I have used BibTeX to cite literature. This method is excellent for automating tedious tasks such as citing papers and generating reference lists based on citation styles, saving time and reducing the likelihood of citation

5. *Markdown, Quarto, and R Markdown*

errors. The literature cited is in a separate file, which can be found on one of my GitHub repositories.

6. Create and host a website

6.1. Creating a website with Quarto

This tutorial guides you through creating a simple, yet professional-looking website using Quarto.

Step W1: Install Quarto

Ensure Quarto is installed on your system. If not, download and install it from [Quarto's official website](#).

Step W2: Create a website

Follow the tutorial that you find [here](#).

Step W3: Copy the `_site` directory

After you have rendered your website a directory “`_site`” appears in the project folder that contains your website. Copy all files of that directory to a directory where you want to save your website. Let's say `my_website`.

In the terminal you can do this with

```
mkdir /home/sthu/my_website/  
cp -r /home/sthu/quarto_website/_site/* /home/sthu/my_website/
```

6.2. Hosting the website on GitHub

R Studio and Quarto offers you various ways to publish the website. I explain you a way that worked out well for me.

Step G1: Create a GitHub account

GitHub will host your thesis website and manage version control for your thesis project. If you don't already have a GitHub account, you'll need to create one: Sign up at [GitHub](#).

Step G2: Create a repository

Create a repository. Name the repo with your username followed by *github.io*. You find a tutorial [here](#).

Step G3: Obtain a personal access token

A personal access token (PAT) is required to authenticate with GitHub from Quarto and RStudio. This token allows you to push changes to your repository securely. Follow the instructions to [create a personal access token on GitHub](#). Alternatively, you can do the following in R:

6. Create and host a website

```
if (!require(pacman)) install.packages("pacman")
pacman::p_load(usethis)
create_github_token()
```

Make sure to note down your token and keep it secure. You'll use this token in RStudio and Quarto to authenticate your GitHub operations.

Step G4: Install and Learn Git

See Section 4.2.

Step G5: Upload the website to GitHub

Use the Terminal of R Studio. Go to the directory with your website that you have copied in Step W3. Then initiate a git repository on the command line, connect it to the repository created in Step G2 on GitHub and finally push it:

```
cd /home/sthu/my_website/
echo "# test" >> README.md
git init
git add README.md
git commit -m "first commit"
git branch -M main
git remote add origin https://github.com/test-hsf/test.git
git push -u origin main
```

Alternatively, you can clone a repository, make some changes, and then push those changes back to GitHub. Here are the Bash commands to accomplish this:

```
# Clone the repository
git clone https://github.com/your-username/your-repository.git

# Make changes, here adding a new file as an example
echo "Some content for the new file" > newfile.txt

# Add the new file to the repository
git add newfile.txt

# Commit the changes
git commit -m "Add new file"

# Push the changes back to GitHub
git push origin main
```

References

- Paul C. Bauer and Camille Landesvatter. Writing a reproducible paper with rstudio and quarto. Technical report, 2023. URL <https://doi.org/10.31219/osf.io/ur4xn>.
- Jennifer Bryan. Happy git and github for the user. URL <https://happygitwithr.com/>.
- Thomas H Davenport and DJ Patil. Data scientist: The sexiest job of the 21st century. *Harvard Business Review*, 90(5):70–76, 2012.
- Christopher Gandrud. *Reproducible research with R and R studio*. Chapman and Hall/CRC, 3 edition, 2020.
- Aud Halbritter and Richard J Telford. Version control with git and github. URL <https://biostats-r.github.io/biostats/github/>.
- Stephan Huber. Quantitative methods, 2024a. URL <https://hubchev.github.io/qm/>.
- Stephan Huber. How to use R for data science, 2024b. URL <https://hubchev.github.io/ds/>.
- Stephan Huber. Managerial economics, 2024c. URL <https://hubchev.github.io/me/>.
- Stephan Huber and Christoph Rust. Calculate travel time and distance with openstreetmap data using the open source routing machine (osrm). *The Stata Journal*, 16(2):416–423, 2016.
- Ben Jones. *Avoiding Data Pitfalls: How to Steer Clear of Common Blunders When Working with Data and Presenting Analysis and Visualizations*. John Wiley & Sons, Hoboken, New Jersey, 2020.
- Roger Martin. How successful leaders think. *Harvard Business Review*, 85(6):71–81, 2007.
- OpenAI. Chatgpt. Large language model, 2024. URL <https://chat.openai.comt>.
- Richard J Telford. Enough markdown to write a thesis, 9 2023. URL <https://biostats-r.github.io/biostats/quarto/>.
- Wikipedia. Survivorship bias, 2024. URL https://en.wikipedia.org/wiki/Survivorship_bias.
- Jeffrey M. Wooldridge. *Introductory Econometrics: A Modern Approach*. South-Western, 2nd edition, 2002.
- Anna C Wysocki, Katherine M Lawson, and Mijke Rhemtulla. Statistical control requires causal justification. *Advances in Methods and Practices in Psychological Science*, 5(2), 2022. URL <https://doi.org/10.1177/25152459221095823>.
- Yihui Xie, Joseph J. Allaire, and Garrett Golemund. *R Markdown: The Definitive Guide*. Chapman and Hall/CRC, 2018.
- Yihui Xie, Christophe Dervieux, and Emily Riederer. *R Markdown Cookbook*. Chapman and Hall/CRC, 2020. available at <https://bookdown.org/yihui/rmarkdown-cookbook>.

A. Exam(ple) question

A.0.1. Question

Please feel free to download a collection of questions that should give you an idea about how I will ask questions in an exam [here](#).

A.0.2. Solutions

Please feel free to download a collection of solutions that should give you an idea about how you can answer the questions in an exam [here](#).