Engineering Statistics

A Handbook of Statistical Inference, Regression, and Experimentation for Engineers

Mansur M. Arief
Industrial & Systems Engineering
King Fahd University of Petroleum & Minerals (KFUPM), Saudi Arabia

ISE 315
2026

Preface

This book is an accompanying material for the course ISE 315, Engineering Statistics, offered by the Industrial & Systems Engineering Department at King Fahd University of Petroleum & Minerals (KFUPM), Saudi Arabia. It has grown into a handbook of statistical inference, regression, and experimentation for engineers, organized in six parts of content plus a worksheets part. Part I builds the foundations: the role of statistics in engineering, how engineering data are collected and summarized, point estimation, and sampling distributions with the central limit theorem. Part II develops inference for a single sample: confidence intervals for means, variances, and proportions, and the full machinery of hypothesis testing. Part III compares two populations: two-sample procedures for means (with variances known, pooled, unequal, and paired), proportions, and variances. Part IV builds regression models: simple linear regression with its inference and diagnostics, then multiple linear regression. Part V treats the design and analysis of experiments: single-factor analysis of variance with multiple comparisons and blocking, and two-factor and factorial experiments. Part VI closes with statistics in practice: a software-and-tools chapter that maps every method in the book to its spreadsheet and Python implementation. A final part collects the printable worksheets, the formula sheet, and the statistical tables used throughout the course.

The reference textbook for the course is Montgomery and Runger, *Applied Statistics and Probability for Engineers*, 6th edition (Wiley). This book follows the notation of that textbook exactly (μ, σ^2, $\bar{X}$, S^2, $\hat{p}$, $z_{\alpha/2}$, $t_{\alpha/2,n-1}$, and so on), so the two can be read side by side, and a margin reference at the start of each section points to the matching textbook section. The book is nevertheless written so that it can be read on its own. You do not need the lecture recordings or the textbook next to you while you read it, although both are useful. Every method

is explained from the beginning, every formula is derived so that you can see where it comes from, and every kind of problem that appears on the homework, the quizzes, and the major exams is solved here in full, step by step.

How this book is organized

Each chapter follows the order of the course lectures. The body of the page carries the explanation, the derivations, and the worked examples. The wide margin carries supporting material, so you always know where to look:

Definitions. Every technical term, when it is first used, is defined in the margin in green.

Figures. Small statistical plots—dot diagrams, distribution curves, table excerpts with the used cell highlighted—sit in the margin, next to the sentence that refers to them, so the picture and the words stay together. Larger panels that you must study closely are placed in the body at full width.

References. Pointers to the matching section of Montgomery and Runger, and to other sources, are placed in the margin at the start of each section.

Lecture pointers. This book is handed out chapter by chapter as the course proceeds, so each chapter tells you exactly which lecture covers what. A green margin note (''Lecture N'') marks the first section of each lecture's coverage, and a green rule across the body with the words ''Lecture N starts here'' marks the exact point where that lecture's material begins. If you read the block between two rules before the corresponding class, you arrive prepared.

Startup lens. Gold margin remarks that reread the section's method from the point of view of someone building a new venture: designing an A/B test, reading churn data, deciding when a pilot result is strong enough to act on.

Engineering judgment. Maroon margin remarks that show where your engineering knowledge—process physics, measurement limits, capacity constraints, failure modes—enters the statistical analysis as an input, not an afterthought.

Safety lens. Maroon margin remarks that connect the method to AI safety testing and to verification and validation practice: how a confidence interval bounds a failure rate, why a test plan for an autonomous system is a sample-size problem, what a p-value does and does not certify.

Numbered procedures

New in this book, every multi-step recipe in the course is set out as a numbered, boxed *Procedure*: constructing a confidence interval, the seven steps of a hypothesis test, fitting a regression line, building an ANOVA table. Procedures are numbered within their chapter (''Procedure 3.2'' and so on), and the lecture slides use the *same* numbers, so when a slide says ''apply Procedure 3.2,'' this book is where that procedure lives, with its derivation and worked examples around it. When you solve a problem, citing the procedure you used is part of showing your work.

Examples, figures, and tables

Every worked *example* is numbered within its chapter (''Example 2.3'' and so on) so it can be referenced in lecture and in your own notes. Inside each example a small *Solution.* rule separates the problem statement from the worked-out solution, leaving you a natural place to attempt the problem yourself before reading on. Every figure and every table is also numbered and listed in the front matter, so the *List of Figures* and the *List of Tables* can be used as a quick visual index of the book. Equations that you will reuse are numbered as well, and the lecture slides cite those equation numbers, so the book and the slides always point to the same formula.

Practice: quizzes, activities, and exercises

Four kinds of practice are built into every chapter:

Pre-class quizzes. short concept checks placed at the start of the chapter, before the first lecture it covers. The questions appear consecutively in a small block; the answers are collected in a shaded box at the bottom of the same block, so you can answer all of them first and then check.

In-class activities. each lecture has a matching in-class activity (ICA), printed at the end of its chapter as a standalone page that can be extracted and used directly in class. ICAs are 15–20 minutes of team work, mostly reasoning rather than computing, with space to write and an exit check at the end.

End-of-chapter exercises. full computational problems in the style of the homework and the quizzes, arranged from basic descriptive calculations up to complete inference problems. Every solution is worked in full, but is collected in the back of the book (the ''Solutions to End-of-Chapter Exercises'' chapter), so that the temptation to glance at the answer while attempting the problem is removed.

Exam-style problems. each chapter's exercises include problems written in the exact format and difficulty of the major exams and the final, with point values marked, so the step from practice to exam holds no surprises.

Each chapter's exercises close with a group called *assessing outside recommendations*: problems in which an AI assistant, a vendor, or a consultant has already produced a statistical analysis—and that analysis is wrong or incomplete in some way. Your task is to find the error, explain it, and correct it. Checking someone else's statistics is a different skill from producing your own, and in practice it is the one you will use most.

Multiple-choice questions, where they appear, list one option per indented line below the stem, so the choices are easy to scan.

A note on numbers and examples

Money amounts, where they appear, are in Saudi riyals. The examples, quizzes, activities, and exercises are built around three recurring settings, chosen so that the same statistical method is seen doing three different jobs: *AI safety testing* (validating a perception model's failure rate, red-team defect discovery, guardrail false-positive rates, autonomous-vehicle disengagement data), *supply chains and transportation* (port container dwell times, truck fleet fuel use, warehouse picking times, delivery lateness), and *building your own startup* (A/B tests on signup conversion, user churn, pilot-customer response times, pricing experiments). Small differences between your answers and the book's usually come from rounding, or from the number of decimal places carried in the statistical tables, and are not errors.

This resource was prepared with the help of generative AI tools and reviewed and corrected by the instructor.

Contents

PART II INFERENCE FOR A SINGLE SAMPLE

PART III COMPARING TWO POPULATIONS

PART IV REGRESSION MODELING

PART VII WORKSHEETS AND APPENDICES

List of Figures

List of Tables