BRIEF CONTENTS

Preface xxi
Acknowledgments xxvii
Introduction xxix

PART I: THE FOUNDATION

Chapter 1: Post-Training Essentials: What It Is and Why It Matters 3
Chapter 2: Prerequisites for Success: Before You Fine-Tune 27

PART II: THE TOOLS

Chapter 3: Supervised Fine-Tuning: The Foundation Technique 67
Chapter 4: Reinforcement Learning: Better Each Time 101
Chapter 5: Preference Optimization: Modern Alternatives to PPO 131
Chapter 6: Evaluation Strategies: Measuring Model Quality 157

PART III: THE CRAFT

Chapter 7: Efficiency Techniques: Quantization and Compression 189
Chapter 8: Domain Adaptation: Make It Yours 221
Chapter 9: Agentic Models: Deeds, Not Words 243
Chapter 10: Reasoning Capabilities: Training for Complex Thought 267

PART IV: THE FRONTIER

Chapter 11: Synthetic Training: Self-Play and Generated Data 295
Chapter 12: Multimodal Systems: Post-Training Beyond Text 319
Chapter 13: Future Directions: What Comes Next 349

Notes 361
Index 375