

# Foundations of Agentic AI for Retail

Concepts, Technologies, and Architectures for Autonomous Retail Systems



Dr. Fatih Nayebe

---

# Foundations of Agentic AI for Retail

Concepts, Technologies, and Architectures for Autonomous Retail Systems

DR. FATIH NAYEBI

Gradient Divergence

Montréal, QC, Canada

# Foundations of Agentic AI for Retail

Concepts, Technologies, and Architectures for Autonomous Retail Systems

Copyright © 2025 Dr. Fatih Nayebi.

**Edition:** First

**Publication Date:** May 5, 2025

**ISBN:** 978-1-0694226-0-6

**Publisher:** Gradient Divergence

**Location:** Montréal, QC, Canada

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means—including photocopying, recording, or other electronic or mechanical methods—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

For permission requests or inquiries, please contact the publisher at:

[contact@gradientdivergence.com](mailto:contact@gradientdivergence.com)

While every precaution has been taken in the preparation of this book, neither the author nor the publisher assumes any liability for errors or omissions, or for damages resulting from the use of the information contained herein.

## Code Repository

All code examples from this book are available in the GitHub repository at <https://github.com/gradient-divergence/agentic-retail-foundations>.

## Community

To join discussions, access additional resources, or participate in Agentic AI projects, visit the Gradient Divergence community at [gradientdivergence.com](https://gradientdivergence.com).

# Epigraph

“The question of whether a computer can think is no more interesting than the question of whether a submarine can swim.” — *Geoffrey Hinton*<sup>1</sup>

“If we’re successful in building truly intelligent systems, we’ll have the biggest opportunity in human history to make the world better for all of humanity. If we fail to build systems aligned with human values, however, we’ll probably have the biggest catastrophe in human history.” — *Stuart Russell*<sup>2</sup>

“The reinforcement learning problem is the AI problem, if you think AI is about an agent. An agent needs to interact with an environment, and learn from its interactions how to improve itself.” — *Richard S. Sutton*<sup>3</sup>

These quotes from AI pioneers frame the profound relationship between artificial intelligence and humanity. They highlight both the immense potential and critical challenges in developing Agentic AI systems that benefit society.

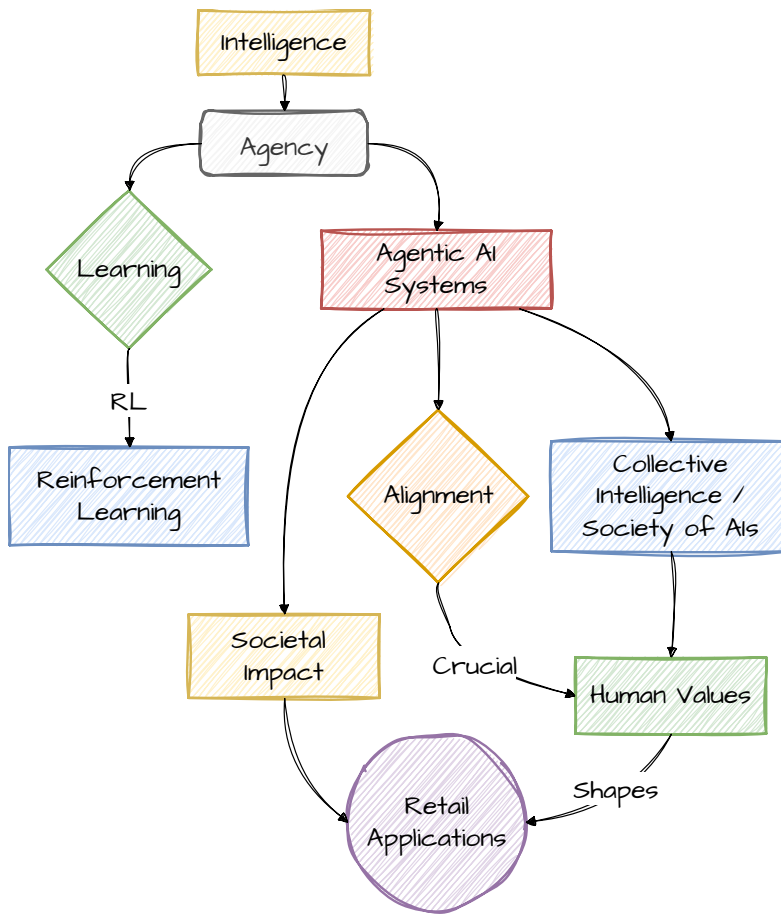
As you explore this book, consider how the foundational principles of Agentic AI must be shaped by human values to create retail systems that augment rather than replace human capabilities.

---

<sup>1</sup>Geoffrey Hinton: Often called a “Godfather of AI,” known for his pioneering work on artificial neural networks and deep learning, particularly backpropagation and Boltzmann machines. Awarded the Turing Award in 2018.

<sup>2</sup>Stuart Russell: Leading AI researcher, co-author of the standard textbook “Artificial Intelligence: A Modern Approach.” Known for his work on rational agents and his advocacy for AI safety and value alignment.

<sup>3</sup>Richard S. Sutton: A key figure in reinforcement learning (RL), co-author of the foundational textbook “Reinforcement Learning: An Introduction.” Known for developing temporal difference learning and actor-critic methods.



### Points to Ponder

- How might Hinton’s analogy about submarines and swimming apply to specific Agentic AI tasks within a retail environment (e.g., inventory management, customer service bots)?
- Considering Russell’s warning, what specific “human values” are most critical to embed in retail AI agents to avoid negative consequences?
- Based on Sutton’s quote, what kinds of “interactions” might a retail agent learn from in a physical store versus an online store?

# Table of Contents

Foreword	xiii
Preface	xv
<b>1 Introduction</b>	<b>1</b>
1.1 From Algorithms to Agents: The Evolution of AI in Retail	3
1.2 What is Agentic AI?	5
1.3 Core Technologies and Architectures Enabling Agentic AI	10
1.4 Key Considerations and Takeaways	17
1.5 Conclusion	19
1.6 Review Questions	20
1.7 Practice Exercises	20
<b>Part I: Foundations of Agentic AI</b>	<b>21</b>
<b>2 Agent Architectures and Frameworks</b>	<b>23</b>
2.1 Defining the Modern AI Agent in Retail	24
2.2 Belief-Desire-Intention (BDI) Models: AI with Human-Like Decision Making	25
2.3 Inside the Mind of a BDI Agent	26
2.4 OODA: Agile Decision Cycles for Dynamic Retail Environments	45
2.5 Choosing the Right Architecture	46
2.6 Bridging Classical Architectures and Modern LLM Patterns	53
2.7 ReAct: Synergizing Reasoning and Acting in LLM Agents	53
2.8 Advanced Agentic Patterns and Frameworks	54
2.9 Choosing the Right Architecture	56
2.10 Conclusion	56
2.11 Review Questions	57
2.12 Practice Exercises	57
<b>3 Decision-Making Frameworks – Statistical &amp; Causal</b>	<b>59</b>
3.1 Decision-Making Process Overview	61
3.2 Optimization Models for Retail Decision-Making	63
3.3 Bayesian Decision Theory	67
3.4 Conclusion	77
3.5 Review Questions	78
3.6 Practice Exercises	78
<b>4 Decision-Making Frameworks – Sequential</b>	<b>79</b>
4.1 Markov Decision Processes (MDPs)	79
4.2 Partially Observable MDPs for Retail Environments	94
4.3 Conclusion	97
4.4 Review Questions	98
4.5 Practice Exercises	98
<b>5 Decision-Making Frameworks – RL &amp; Planning</b>	<b>99</b>
5.1 Reinforcement Learning: Learning Through Interaction	99

5.2	Planning and Optimization in Retail Decisions	107
5.3	Code Example: Store Fulfillment Optimization	110
5.4	Conclusion	120
5.5	Review Questions	121
5.6	Practice Exercises	121
<b>Part II: Enabling Technologies and Architectures</b>		<b>123</b>
<b>6</b>	<b>Foundation Models and Visual Intelligence</b>	<b>125</b>
6.1	Critical Technological Pillars	126
6.2	Large Language Models as Reasoning Engines	131
6.3	Computer Vision for Physical Store Awareness	143
6.4	Conclusion	155
6.5	Review Questions	156
6.6	Practice Exercises	156
<b>7</b>	<b>Sensor Networks and Cognitive Systems</b>	<b>157</b>
7.1	IoT and Sensor Networks: The Nervous System of Retail Agents	159
7.2	Knowledge Graphs and Semantic Reasoning: Structuring Retail Intelligence	171
7.3	Causal Reasoning and Counterfactual Analysis in Retail	185
7.4	Conclusion	199
7.5	Review Questions	200
7.6	Practice Exercises	200
<b>Part III: Multi-Agent Systems and Integration</b>		<b>201</b>
<b>8</b>	<b>Multi Agent Systems in Retail</b>	<b>203</b>
8.1	Why Multi-Agent Systems for Retail?	203
8.2	Understanding Multi-Agent Systems (MAS) in Retail	204
8.3	Applications of Multi-Agent Systems in Retail Operations	207
8.4	Agent Communication Protocols in Retail	208
8.5	Multi-Agent System Architectures in Retail	210
8.6	Collaborative vs. Competitive Agent Systems	230
8.7	Conclusion	237
8.8	Review Questions	238
8.9	Practice Exercises	238
<b>9</b>	<b>End-to-End Integration for Autonomous Retail</b>	<b>239</b>
9.1	System Architecture Overview	240
9.2	Core Principles for End-to-End Integration	242
9.3	The Integration Journey	242
9.4	Agent Workflow Management	243
9.5	Event-Driven Architectures	254
9.6	API-Based Communication Between Agents	264
9.7	State Management Across Agent Systems	272
9.8	Human Interaction in Multi-Agent Systems	279
9.9	Real-Time Decision Making and Feedback Loops	281
9.10	Conclusion	288
9.11	Review Questions	289
9.12	Practice Exercises	290
<b>Part IV: Implementation and Ethical Considerations</b>		<b>291</b>
<b>10</b>	<b>Implementing Agentic Systems in Retail</b>	<b>293</b>



10.1	Implementation Workflow . . . . .	294
10.2	Agent Development Methodologies . . . . .	299
10.3	Monitoring and Maintaining Agent Systems . . . . .	307
10.4	Enterprise Scaling Challenges for Retail Agents . . . . .	316
10.5	Conclusion . . . . .	317
10.6	Review Questions . . . . .	318
10.7	Practice Exercises . . . . .	318
<b>11</b>	<b>Operational Excellence for AI Engineering in Retail</b>	<b>319</b>
11.1	DevOps Foundations: From Commit to Running Container . . . . .	319
11.2	Workflow Engines for Complex Retail Processes . . . . .	322
11.3	Observability: Seeing the Whole Elephant . . . . .	323
11.4	The Interconnected Lifecycle: DevOps, DataOps, MLOps . . . . .	324
11.5	DataOps: Trustworthy Pipelines . . . . .	324
11.6	MLOps Lifecycle . . . . .	325
11.7	Continuous Evaluation & Experimentation . . . . .	326
11.8	CI/CD Pipeline Blueprint . . . . .	326
11.9	Operational KPIs . . . . .	329
11.10	Case-in-Point: FastAPI Latency Middleware . . . . .	330
11.11	Security & Compliance: Protecting Customer Trust . . . . .	330
11.12	Incident Response & Chaos Engineering . . . . .	331
11.13	Cost Optimisation & FinOps for Agents . . . . .	331
11.14	Infrastructure as Code & Platform Engineering . . . . .	332
11.15	Advanced GitOps Patterns . . . . .	332
11.16	SRE Playbooks & On-call Excellence . . . . .	333
11.17	Case-in-Point: Global Retailer Black Friday . . . . .	334
11.18	Future Trends & Emerging Tools . . . . .	334
11.19	Self-Audit Checklist . . . . .	335
11.20	Conclusion . . . . .	335
11.21	Review Questions . . . . .	336
11.22	Practice Exercises . . . . .	336
<b>12</b>	<b>Ethical Considerations and Governance</b>	<b>337</b>
12.1	Ethical Governance Framework . . . . .	337
12.2	Transparency and Explainability . . . . .	339
12.3	Accountability for Agent Decisions . . . . .	342
12.4	Human-in-the-Loop Approaches . . . . .	348
12.5	Risk Management for Autonomous Systems . . . . .	356
12.6	Conclusion . . . . .	363
12.7	Review Questions . . . . .	364
12.8	Practice Exercises . . . . .	364
	<b>Part V: Case Studies and Future Directions</b>	<b>365</b>
<b>13</b>	<b>Real-World Case Studies</b>	<b>367</b>
13.1	Autonomous Inventory Management . . . . .	368
13.2	Agentic Pricing and Promotion Systems . . . . .	375
13.3	Customer-Facing Retail Agents . . . . .	381
13.4	Conclusion . . . . .	388
13.5	Review Questions . . . . .	389



13.6 Practice Exercises . . . . .	389
<b>14 Summary and Future Directions</b>	<b>391</b>
14.1 Key Takeaways for Retail Implementers . . . . .	391
14.2 Emerging Trends in Agentic Retail . . . . .	395
14.3 The Path to Fully Autonomous Retail . . . . .	397
14.4 Conclusion: Charting the Course for Agentic Retail . . . . .	401
14.5 Review Questions . . . . .	402
14.6 Practice Exercises . . . . .	402
<b>Appendix A: Advanced Mathematical Foundations for Decision Frameworks</b>	<b>403</b>
<b>References</b>	<b>409</b>
<b>About the Author</b>	<b>413</b>

# List of Figures

1.1	AI adoption in retail	2
1.2	Benefits of AI adoption	2
1.3	Evolution of AI in Retail	3
1.4	Evolution of AI in retail	4
1.5	Agentic AI capabilities	4
1.6	Agentic Retail System Architecture	6
1.7	Agentic AI System Loop	7
1.8	Example of interaction between customers, Agentic AI, data systems, and external services	9
1.9	Layered Reference Architecture with Cross-Cutting Concerns	14
2.1	Belief-Desire-Intention (BDI) Architecture	25
2.2	Agent Data Models and Architecture	29
2.3	OODA Loop	45
2.4	OODA-Based Dynamic Pricing Agent	48
3.1	Decision Making Frameworks	60
3.2	Decision Making Framework Selection Approach	61
3.3	Decision Making Process	62
3.4	Structured Approach to Decision Making Process	62
3.5	Bayesian Decision Theory	67
3.6	Bayesian Product Recommendation Agent	70
4.1	Markov Decision Process	80
4.2	Example: MDP for Dynamic Pricing	87
5.1	Reinforcement Learning Cycle	100
5.2	Store Fulfillment Optimization	110
6.1	Core Technologies Integration	126
6.2	Critical Technological Pillars	127
6.3	Typical LLM Integration Workflow	128
6.4	Integrated Agentic System Addressing an Inventory Shortage	131
6.5	LLM as Reasoning Engine	132
6.6	LLM-Powered Customer Service Agent	136
6.7	Computer Vision for Store Awareness	143
6.8	Computer Vision for Shelf Monitoring Architecture	148
7.1	Sensor Network Architecture	158
7.2	Retail Sensor Network	159
7.3	Privacy-Preserving Edge Gateway	163
7.4	Processing Sensor Data for Real-Time Agent Decisions	164
7.5	Retail Knowledge Graph	172
7.6	Knowledge Graph for Retail Product Relationships	176
7.7	Causal Inference in Retail	185
7.8	Causal Inference for Promotion Effectiveness	188
7.9	Causal Reasoning integrating with other agent systems	198

8.1	Multi-Agent Systems in Retail Operations . . . . .	207
8.2	FIPA Message Example . . . . .	209
8.3	Coordination between different agents and external systems in a retail environment . . . . .	211
8.4	Detailed Multi-Agent Retail System Architecture . . . . .	211
8.5	Implementation of Agent Communication . . . . .	216
8.6	MAS Coordination Mechanisms (Centralized vs. Decentralized) . . . . .	219
8.7	CNP for Task Allocation . . . . .	220
8.8	Task Allocation Among Store Agents . . . . .	221
8.9	Auction Mechanism for Supplier Selection . . . . .	226
8.10	Collaborative vs. Competitive Agent Systems . . . . .	230
8.11	Cooperative Inventory Sharing Between Stores . . . . .	232
9.1	End-to-End Retail Integration Architecture . . . . .	240
9.2	The Integration Journey . . . . .	243
9.3	CQRS Pattern in Retail . . . . .	255
10.1	Reference Architecture for a Retail Agentic System (cloud & Edge) . . . . .	298
11.1	The Interconnected Lifecycle: DevOps, DataOps, MLOps . . . . .	324
11.2	CI/CD Pipeline Blueprint . . . . .	326
12.1	Ethical Governance Framework . . . . .	338
12.2	Governance workflow for an AI agent from design to monitoring . . . . .	347
12.3	Workflow for an escalation scenario . . . . .	352
13.1	Inventory Management Cycle . . . . .	368
13.2	A typical multi-agent workflow for autonomous inventory . . . . .	370
13.3	Agents in a dynamic pricing system and their interactions . . . . .	376
13.4	Interaction of a customer-facing chatbot agent with backend systems to personalize responses . . . . .	383
14.1	Maturity Model for Agentic Retail . . . . .	394
14.2	Structured timeline for anticipated advancements . . . . .	400

# Foreword

By Professor Alain Abran, Ph.D., Ing.

Emeritus Professor, Department of Software Engineering and IT

École de technologie supérieure (ÉTS), Montréal

---

When I first met Fatih as a doctoral candidate in software engineering, his curiosity was already leaning toward the then-nascent field of machine learning. Back then, discussions of autonomous agents and large-scale AI systems were still largely confined to research seminars and speculative conferences; few imagined the sweeping industrial impact we witness today. Yet Fatih was convinced—even then—that rigorous engineering principles could (and should) underpin intelligent systems long before “AI” became a ubiquitous business acronym.

Over the years we spent together—first during his Ph.D., co-supervised with my colleague Jean-Marc Desharnais, and later while he served as a post-doctoral researcher in our laboratory—we co-authored publications that blended empirical measurement with innovative uses of predictive models. Those collaborations affirmed a shared conviction: software engineering, when anchored in disciplined methods and robust bodies of knowledge, can adapt and thrive even as the underlying technologies evolve at breakneck pace.

That conviction lies at the heart of *Foundations of Agentic AI for Retail*. The book you are about to read is not merely a technical manual, though it abounds in architectural blueprints, code examples, and implementation guides. Nor is it purely an industry playbook, though retail leaders will find it invaluable for translating AI hype into operational advantage. It is, instead, a bridge—between scientific rigor and real-world applicability, between the enduring principles codified in the *SWEBOK* and the frontier concepts now reshaping commerce through autonomous agents.

## A rigorous lineage

In my own career, I have argued that software engineering must remain rooted in measurable evidence and systematic knowledge. The *Software Engineering Body of Knowledge (SWEBOK)* was conceived to provide practitioners with a stable, shared foundation—much as civil engineers rely on structural mechanics or physicians on anatomy. Fatih extends that philosophy into the realm of Agentic AI. From his lucid treatment of Belief-Desire-Intention (BDI) models and OODA loops, to his detailed guidance on reinforcement learning pipelines and event-driven architectures, he demonstrates that even the most sophisticated AI agents can—and must—be engineered with the same care we devote to any critical system.

## Why retail, why now?

Retail may seem, at first glance, an unlikely vanguard for Agentic AI. Yet few industries present a richer tapestry of real-time signals—prices, inventories, customer behaviors, supply-chain events—demanding rapid, decentralized decisions. Fatih’s choice of retail as a proving ground is therefore inspired: it exposes every limitation of monolithic, rule-based software and makes a compelling case for autonomous, collaborative agents governed by clear objectives, guardrails, and feedback loops.

Readers will appreciate how seamlessly the book weaves advanced theory with concrete practice. Chapter-by-chapter, Fatih moves from foundational concepts to decision-making frameworks, enabling technologies, multi-agent coordination, and finally to full end-to-end integration—including the

ethical and governance considerations that responsible engineers must never overlook. The result is a text that will guide C-suite executives, software architects, data scientists, and graduate students alike.

### The human dimension

Underlying the algorithms and patterns is Fatih's conviction that technology ultimately serves human progress. His emphasis on Human-in-the-Loop safeguards, transparency, and rigorous evaluation echoes the broader movement toward responsible AI—an ethos that aligns with the scientific mindset we fostered at ÉTS. I am particularly pleased to see extensive attention given to explainability, accountability, and risk management, ensuring that Agentic AI advances do not outpace our capacity to govern them.

### A glance toward the horizon

Agentic systems will soon permeate domains far beyond retail—healthcare, energy, transportation, public services—wherever complex, dynamic environments require continuous adaptation. The frameworks articulated here will serve as a template for those future applications. More importantly, they remind us that even as AI models grow in capability, the disciplines of requirements engineering, measurement, validation, and ethical oversight remain indispensable.

Fatih has delivered a timely, authoritative, and engaging work. It is a testament to his evolution from inquisitive graduate student to industry leader and educator, and it reflects the very principles we strived to instill: intellectual curiosity, methodological rigor, and an unwavering focus on practical impact.

I invite you, the reader, to dig into these pages with both critical attention and creative imagination. May you emerge not only informed but inspired to engineer the next generation of intelligent systems—systems that honor the best traditions of our discipline while venturing boldly into new frontiers.

*Montréal, April 2025*

Alain Abran

# Preface

## A Meeting of Theory and Practice

The retail industry is in a period of unprecedented upheaval, driven by rapid advances in technology and seismic shifts in consumer behavior. As artificial intelligence (*AI*) emerges from research labs and enters the mainstream, retailers grapple with a wave of new possibilities—smart shelves that reorder themselves, personalized promotions that adapt in real time, and automated systems that anticipate trends before they become trends. Yet, for every promising pilot project, there remains a wide chasm between conceptual experimentation and fully realized, at-scale *Agentic AI* solutions.

Over the years, I have observed this tension from two vantage points: the technology sector, where startups and established companies alike innovate at breakneck speed, and the academic world, which rigorously interrogates the underlying theory and ethics of *AI*. In both spheres, the concept of the “*autonomous agent*”—a software entity capable of perceiving its environment, reasoning about complex states, and taking decisive action—has sparked keen interest. But while the term “*Agentic AI*” has found its way into research papers and conference keynotes, the practical guidance for deploying such systems in the dynamic realm of retail remains sparse.

## Why Now?

We stand at a pivotal moment. The retail industry faces surging expectations from consumers who demand instant gratification, endless customization, and seamless offline-to-online experiences. Traditional methods—largely reliant on human-driven decision-making and *heuristic-based approaches*—are buckling under the weight of these expectations. Meanwhile, *AI*-driven breakthroughs in *computer vision*, *natural language processing*, *reinforcement learning*, and *edge computing* have given us the technical tools needed to build more adaptive and self-sufficient systems.

These converging forces have created an urgent need for a unifying, accessible resource that synthesizes the full range of *Agentic AI* capabilities, from foundational theories to architectural best practices. This book aims to fill that void, offering a step-by-step journey through the fundamentals of *agent design*, *decision frameworks*, *multi-agent coordination*, and *end-to-end integrations* for real-world retail contexts.

## Who This Book Is For

### Quick Guide: What’s In It For You?

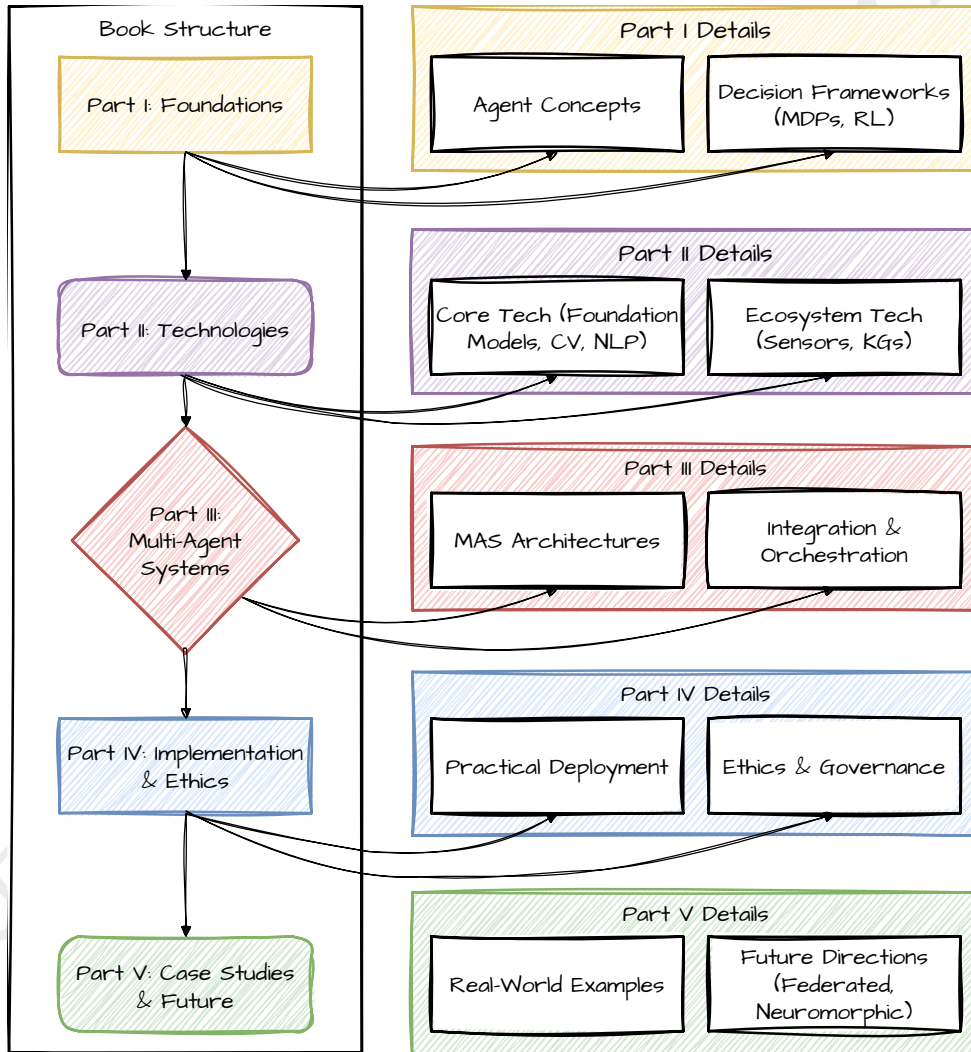
- **Executives:** Understand strategic value, applications (supply chain, CX), and implementation success factors for *Agentic AI*.
- **Engineers/Scientists:** Gain practical architectural insights, explore libraries/code examples, and bridge theory with production-grade *AI*.
- **Product Managers/Analysts:** Grasp the “why” and “how” of *agentic systems* to align stakeholders and technical feasibility.
- **Academics/Instructors:** Find real-world retail *AI* case studies and deployment examples to connect research to practice.

1. **Retail Executives and Decision-Makers** If your role involves strategic planning or high-level oversight, you’ll find clarity here on how *Agentic AI* can reshape key areas of retail—supply chain optimization, customer experience, and more—while uncovering common pitfalls and strategies for success.

2. **Data Scientists and Engineers** Technical teams charged with creating or maintaining AI-driven solutions will gain practical insights into architectures, libraries, and coding examples. Think of this as your guide for bridging theoretical AI algorithms with robust, production-grade implementations.
3. **Product Managers and Business Analysts** As the conduit between technical teams and executive leadership, you need a solid grasp of both the “why” and “how” of deploying agentic systems. This book offers a detailed roadmap that will help align stakeholder objectives with technical feasibility.
4. **Academic Researchers and Instructors** Those teaching or researching AI, multi-agent systems, or retail innovation will find real-world case studies illustrating how Agentic AI moves from whiteboard concepts to in-store deployments.

## Scope and Structure

A roadmap from first principles to full-scale deployment





The book is organised in five deliberate movements. Each Part builds on the previous one: first clarifying *what* Agentic AI is, then *how* to build it, *how* to network many agents together, *how* to harden the solution for production, and finally *where* all this is heading. Skim linearly for a masterclass, or jump straight to the Part that solves today’s problem.

Part	Chapters	Core Question	Value Promise	Key Takeaways
I – Foundations of Agentic AI	1 – 5	<i>What makes an agent “agentic”?</i>	Establishes the mathematical and conceptual bedrock—BDI, OODA, Bayesian & causal decision models, MDPs, RL, planning.	Readers leave with a rigorous mental model and reference code for single-agent intelligence.
II – Enabling Technologies & Architectures	6 – 7	<i>Which technologies turn theory into capability?</i>	Dissects LLMs, vision, sensor fabrics, knowledge graphs, causal engines, and their orchestration inside retail platforms.	Blueprint-level diagrams show how to wire perception, reasoning and action into a cohesive stack.
III – Multi-Agent Systems & Integration	8 – 9	<i>How do many agents collaborate (or compete) at retail scale?</i>	Covers MAS topologies, communication protocols (FIPA, MCP, A2A), negotiation, task-allocation patterns, and end-to-end orchestration.	Practical code and patterns for stitching agents across supply-chain, stores, e-commerce and HQ.
IV – Implementation & Ethical Guardrails	10 – 12	<i>How do we ship safely, securely and at enterprise scale?</i>	Walks through Dev/Data/MLOps, observability, CI/CD, SRE, privacy, risk, explainability, and regulatory compliance.	Templates and checklists ensure production readiness and responsible governance from day one.
V – Case Studies & Future Directions	13 – 14	<i>What’s working now, and what’s next?</i>	Deep dives into live deployments—inventory, dynamic pricing, customer agents—and surveys federated learning, neuromorphic & quantum horizons.	Lessons learned, ROI metrics, and a foresight timeline arm readers for the next decade.

## A Collaborative Lens on Agentic AI

This book is the product of many minds—retail operators, data scientists, ethicists, supply-chain strategists, software engineers, and academic researchers—who stress-tested every chapter. Their cross-disciplinary feedback keeps the material clear whether you care about GPU latency, inventory turns, or governance policy. Agentic AI can only reach its full potential when diverse perspectives work in concert; that principle guided every page that follows.

### Reading Paths: Find the Chapters That Serve You Best

- **Executives & Business Leaders (CEO / CMO / COO)**  
Skim the opening section of each chapter for high-level concepts, business impact, and strategic takeaways. Zero-in on Introduction (Ch 1), Implementation Strategy (Ch 10), Ethical Considerations (Ch 12), Case Studies (Ch 13), and Future Directions (Ch 14). The *Key Takeaways* boxes distill the essence without deep technical detail.
- **Architects & Technical Leaders (CTO / Enterprise Architects)**  
After each chapter’s intro, dive into Agent Architectures (Ch 2), Decision Frameworks (Ch 3-5), Core Technologies (Ch 6-7), Multi-Agent Systems (Ch 8-9), and Implementation Workflows (Ch 10). Pay special attention to system diagrams, integration patterns, and *Limitations & Challenges* call-outs to pre-empt real-world hurdles.

- **Mathematicians & Researchers**  
Focus on the formal treatments in **Chapters 2-7** and **Appendix A**. These cover mathematical foundations, proofs, and guarantees that link retail applications to rigorous theory. The extensive *References* section will steer further scholarship.
- **Engineers & Developers**  
Head straight for the hands-on material in **Chapters 2-10**. Complete, runnable code listings, framework walk-throughs, and MLOps blueprints provide everything you need to build, test, and ship agentic systems.

Each chapter follows a consistent arc—**Business Context** -> **Theory** -> **Hands-on Implementation** -> **Key Takeaways**—so you can choose your depth of engagement and still stay on the narrative rail.

## My Journey and Aspirations

My path to writing *Foundations of Agentic AI for Retail* has been shaped by a career spent at the crossroads of enterprise technology, academic research, and practical product development. As Head of Data, Analytics, and *AI* at a global retailer, I have navigated large-scale deployment challenges, from securing organizational buy-in to wrestling with integration complexities. As a Faculty Lecturer, I have found joy in making advanced *AI* concepts accessible to students and professionals who arrive with diverse backgrounds yet share a zeal for innovation.

This book is both a testament to the road traveled and a roadmap for the journey yet to come. My hope is that these pages demystify Agentic AI and act as a catalyst—moving you from proofs-of-concept to production, from tactical wins to strategic transformation. Done well, autonomous agents don't replace humans; they free us to focus on creativity and strategy.

Above all, I hope that by blending practical guidance with deep theoretical underpinnings, *Foundations of Agentic AI for Retail* can be the catalyst that propels you from proofs-of-concept to transformative, industry-leading solutions. The future of retail, I believe, rests on the shoulders of *autonomous agents* that complement human expertise rather than substitute it—creating a world where intelligent systems augment, rather than eclipse, our innate potential.

## Code Repository and Interactive Notebooks

All code examples from this book are available in the GitHub repository at <https://github.com/gradient-divergence/agentic-retail-foundations>. The repository includes *marimo* notebooks for each chapter, allowing you to interact with the code, modify parameters, and experiment with the concepts in real-time. While the code examples presented in the book chapters are designed for clarity and brevity, providing illustrative snippets of core concepts, the repository contains the complete, executable *Marimo* notebooks with more extensive implementations, detailed data handling, and additional features suitable for deeper exploration and experimentation. These interactive notebooks make it easier to understand complex algorithms and see how different parameters affect outcomes in retail-specific contexts.

## Join the Gradient Divergence Community

*Agentic AI* for retail is a rapidly evolving field, and ongoing collaboration is essential for continuing innovation. I invite you to join the Gradient Divergence community at [gradientdivergence.com](https://gradientdivergence.com), where you'll find:

- Regular blog posts on the latest *Agentic AI* developments
- A forum for discussing implementation challenges and solutions
- Access to additional code examples and extended case studies
- Opportunities to connect with other retail technologists and *AI* practitioners

The community is committed to advancing the practical application of *AI* in retail environments and welcomes contributions from practitioners at all levels of expertise.

## Acknowledgments

The journey to create *Foundations of Agentic AI for Retail* has been one of exploration and collaboration, made possible by an extraordinary academic and professional network. I have had the privilege of interacting with thought leaders, students, and practitioners who have shaped my understanding of *Agentic AI* and its implications for retail.

### A Community of Scholars and Innovators

I first thank the faculty and research staff at [McGill University](#), especially within the Desautels Faculty of Management, for fostering a rigorous and intellectually stimulating environment. Their open forums, reading groups, and joint projects challenged and refined my thinking on autonomous systems. I am particularly grateful for the interdisciplinary collaborations that offered diverse perspectives on AI's role in retail.

### Students: The Lifeblood of Inspiration

To the graduate and undergraduate students I've encountered: your curiosity and tenacity in courses like *Enterprise Data Science: Concepts and Algorithms*, *Enterprise Machine Learning in Production*, *Introduction to AI and Deep Learning*, *Applications and Architectures of Deep Learning*, and *Designing and Developing Agentic AI Systems*, as well as in hackathons and seminars, constantly inspired me. Your questions spurred me to re-examine assumptions and seek better solutions. This book greatly benefited from our dialogues.

I also acknowledge the *Retail Gen AI Hackathon* and *Capstone project* teams at [McGill University](#). Your passion for applying theory to practice validated the potential for academia to drive impactful industry solutions, informing many *use cases* and *architectural frameworks* herein.

### Industry–Academia Synergy

My experiences at the [ALDO Group](#) highlighted the power of applied AI. Collaborating with *data scientists*, *engineers*, and *strategists* provided invaluable insights into deploying *Agentic AI* in retail. This manuscript is enriched by the dynamic exchange between academic theory and industry practice. Special thanks to the Data, Analytics, and AI team for their exploratory spirit and feedback.

Gratitude extends to the broader ecosystem of *industry partners*, *research consortiums*, and *AI conferences*. Their collective experiences advanced the field and shaped the *code snippets*, *decision frameworks*, and *multi-agent coordination* strategies presented.

### Technical Reviewers and Early Readers

This book has been strengthened by the critical eyes of the *technical reviewers* and *early readers* who generously devoted their time to dissecting initial drafts. Their rigorous attention to detail, pointed questions, and calls for clarity significantly elevated this work. Contributions from experts across *AI research*, *cloud architecture*, and *large-scale retail systems* helped refine the technical accuracy and contextual relevancy of each chapter.

Notable contributions include:

- **Arial Huang:** Review and insightful distinctions between traditional AI and *Agentic AI*, sharpening the narrative.
- **Armen Momejian:** Provided valuable feedback on book structure and organization as well as insightful suggestions on multiple chapters.

- **Arthur Pentecoste:** Delivered meticulous chapter-by-chapter reviews, identifying areas for improved narrative flow, context, and technical accuracy, including LaTeX corrections.
- **Basant Mounir:** Offered key insights on overall structure and chapter organization, along with helpful feedback on several sections.
- **Chiara Liu:** Offered insights on governance, practical code implementation suggestions, and advanced LLM techniques, enhancing the book's technical depth and usability.
- **Joseph and Roonie Corera:** Provided helpful general feedback, contributing to the overall refinement of the manuscript.
- **Laurence Audrey Vincent:** Detailed feedback on the chapter covering *Ethical Considerations and Governance*, adding essential nuance and depth.
- **Matthieu Houle:** Provided comprehensive feedback across multiple chapters, focusing on conceptual clarity, the integration of scientific approaches in retail operations, and specific figure/example improvements.
- **Necmiye Genc:** Thorough review and thoughtful commentary, offering a fresh lens on structure and substance.
- **Onur Erkin Sucu:** Careful review and invaluable feedback on mathematical components, source code, enhancing clarity and precision.
- **Yael Kochman:** Provided valuable feedback on clarity, structure, and the accessibility of introductory concepts for a broad audience.
- **Yash Joshi:** Contributed valuable suggestions on incorporating recent agent architectures, frameworks, deployment patterns, and industry case studies, ensuring the book's contemporary relevance.

A special mention goes to the reviewers who stress-tested the ideas herein against real-world scenarios. Your unique vantage point—situated at the intersection of academic experimentation and brick-and-mortar realities—offered a grounded perspective that kept the text both forward-thinking and pragmatically sound.

## Looking Ahead

I view this book not as a static endpoint but as part of a living conversation about the evolution of AI in retail. The success of Agentic AI systems depends on open idea exchange, interdisciplinary research, and inclusive dialogue. It is my sincere hope that readers will take these concepts, challenge them, refine them, and push them to new frontiers.

To all those—students, researchers, industry colleagues, and academic peers—who have fueled my passion for teaching and learning, thank you. Your collective contributions have guided me in weaving together the theoretical and practical dimensions of *Agentic AI*. I am humbled by your support and invigorated by the knowledge that together, we stand at the cusp of a transformative era in retail.

— *Dr. Fatih Nayebe*

Montréal, 2025

# CHAPTER 1: Introduction

In this chapter, we explore what makes AI “agentic,” transitioning from traditional methods to autonomous decision-making systems. We’ll discuss foundational concepts, the AI lifecycle, and the essential building blocks that position Agentic AI as a transformative force in retail, enabling a more scientific approach to daily operations. Readers will gain clarity on how proactive intelligence reshapes inventory management, pricing, and customer experiences, setting the stage for deeper exploration in subsequent chapters.

## Learning Objectives

By the end of this chapter, you will be able to:

1. **Conceptual Understanding**
  - Understand the fundamental principles of Agentic AI and its role in retail
  - Recognize the key differences between Agentic AI and traditional AI approaches
  - Identify the core components of Agentic AI systems
2. **Technical Proficiency**
  - Comprehend the sense-think-act loop in Agentic AI systems
  - Understand the architecture of Agentic retail systems
  - Recognize the technical requirements for implementing Agentic AI
3. **Practical Application**
  - Evaluate potential use cases for Agentic AI in retail
  - Assess the benefits and challenges of implementing Agentic AI
  - Understand how Agentic AI can transform retail operations

Retail is at a turning point unlike any we’ve seen before—one defined by the power of **Artificial Intelligence (AI)**. Imagine retailers so agile they can predict customer needs before customers themselves are even aware. Envision intelligent systems autonomously making complex decisions around the clock, from setting prices and optimizing inventory to personalizing customer experiences and anticipating upcoming trends. This isn’t speculative futurism; it’s happening right now, bringing a new level of scientific rigor to retail operations.

AI’s impact on retail strategy is profound, and companies that embrace it thrive while those that hesitate risk obsolescence. Consider the numbers: **87% of retailers have already implemented AI in at least one aspect of their operations**, and **60% plan to make substantial new investments in the near future**. By 2025, **80% of retail executives expect to see wide-scale automation powered by AI in their organizations**—transformations that have already boosted annual revenues for **69% of adopters** and cut operational costs for **72%** (Neontri 2023).

In other words, AI is no longer just an option; it’s the new frontier for retailers determined to remain competitive. Those who harness its capabilities will lead the way, redefining what modern retail can be. Those who don’t will inevitably be left behind. The choice is clear, and the future starts now.

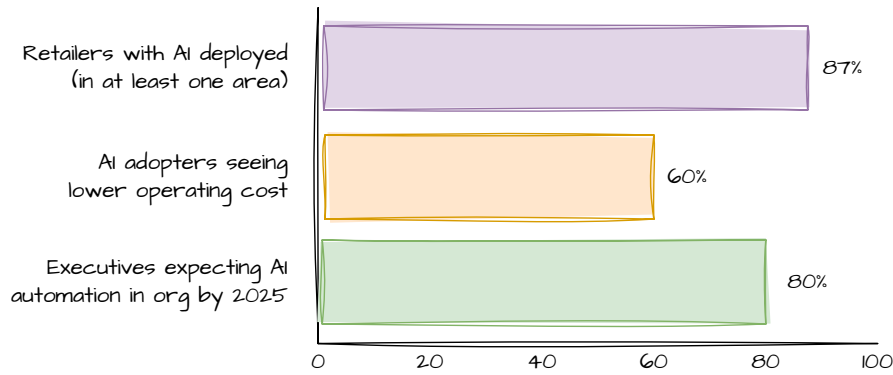


Figure 1.1: AI adoption in retail

*AI adoption in retail is accelerating. A majority of retailers already employ AI in various capacities, with many planning further investments. Executives anticipate broader adoption of AI-driven automation (Neontri 2023).*




Figure 1.2: Benefits of AI adoption

*Retailers are seeing clear benefits from AI. Surveys reveal that 69% report higher annual revenue, and 72% experience lower operating costs, highlighting AI's positive impact.*

#### ! Key Statistics on AI in Retail

- 87% of retailers have deployed AI in at least one area
- 60% plan to significantly boost investments
- By 2025, 80% of retail executives anticipate extensive AI-driven automation
- 69% report increased annual revenues
- 72% have reduced operating costs through AI

Over the past decade, retail AI applications have evolved significantly—from basic analytics and rule-based automation to sophisticated *generative AI* capable of creating content such as product descriptions, personalized recommendations, and customer communications. Today, however, an entirely new frontier has emerged: **Agentic AI** (Wooldridge and Jennings 1995; Brown et al. 2020).

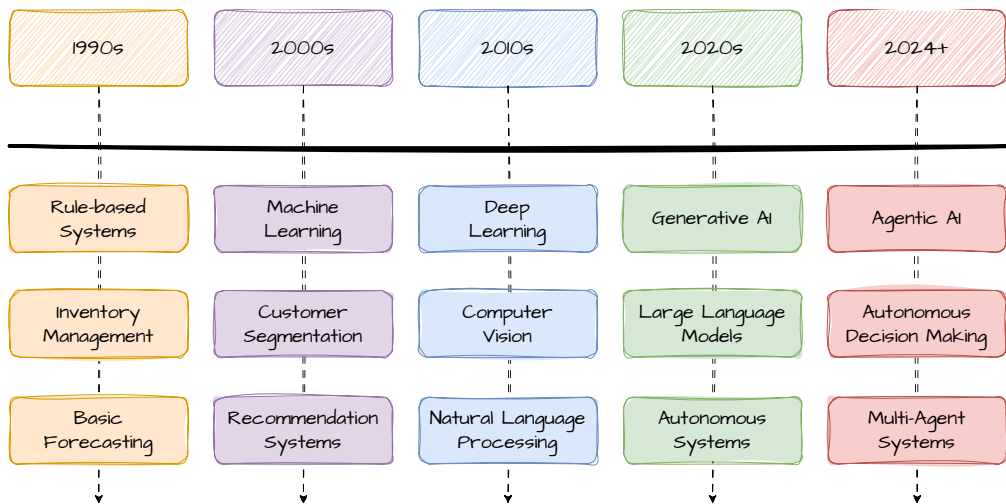


Figure 1.3: Evolution of AI in Retail

Agentic AI brings together the versatility of large language models (LLMs) with the structured decision-making of traditional software, enabling AI systems to not only analyze or generate information, but to **take autonomous actions** in pursuit of goals (IBM Insights 2023; Hitzler, Sarker, and Krisnadhi 2022). In essence, Agentic AI is *proactive* where generative AI is reactive. Rather than waiting for a human prompt at each step, an Agentic AI system can independently decide *what* needs to be done next. This promises to revolutionize retail through autonomous decision-making capabilities that far exceed those of earlier AI systems (Marr 2023).

Imagine walking into a retail store or browsing an e-commerce site where every interaction feels uniquely tailored to you—where systems don’t simply respond to your actions but anticipate your needs, seamlessly adapting to every subtle shift in context. This is no longer the realm of futuristic speculation but the reality of **Agentic AI for Retail**, a groundbreaking approach transforming passive computational tools into autonomous agents that sense their environment, reason about complex scenarios, and proactively take actions aligned with overarching business goals.

Traditional retail technology typically follows rigid, pre-defined instructions, lacking the flexibility to adjust to unpredictable market fluctuations or evolving consumer preferences. Agentic AI transcends these limitations, shifting from static predictive engines toward dynamic, strategic entities. These autonomous agents not only anticipate and plan but also learn from outcomes and improve over time without continuous human intervention. The transformation from reactive systems to proactive, intelligent partners signals a profound evolutionary leap in retail technology, redefining every touchpoint in the customer journey (including Awareness, Consideration, Purchase, Service, and Loyalty stages, covering aspects like Marketing, Advertising, Sales, and Support) and reshaping entire business processes.

This book is your comprehensive guide to understanding, implementing, and leveraging Agentic AI—transforming conventional retail technology from static, responsive tools into dynamic, autonomous strategic partners. Welcome to the future of retail.

## 1.1 From Algorithms to Agents: The Evolution of AI in Retail

The evolution of AI in retail can be viewed in three distinct waves.





Figure 1.4: Evolution of AI in retail

### First Wave: Automation of Routine Tasks

Initially, retail technology was predominantly transactional, focused on automating repetitive tasks such as inventory management, point-of-sale transactions, and basic data processing. These systems, though beneficial, were limited and required constant human oversight and manual intervention.

### Second Wave: Predictive Intelligence through Machine Learning

The introduction of machine learning marked a significant progression, allowing systems to identify patterns and make predictive forecasts. Retailers began utilizing these capabilities for demand forecasting, personalized customer recommendations, and pricing optimization. Despite this sophistication, these technologies remained reactive and were confined within narrow functional silos. They were unable to autonomously adapt to novel scenarios or coordinate across different functions without extensive human reprogramming.

### Third Wave: Emergence of Agentic AI

Agentic AI represents a revolutionary leap forward. These advanced systems exhibit four critical capabilities that distinguish them from earlier AI paradigms.

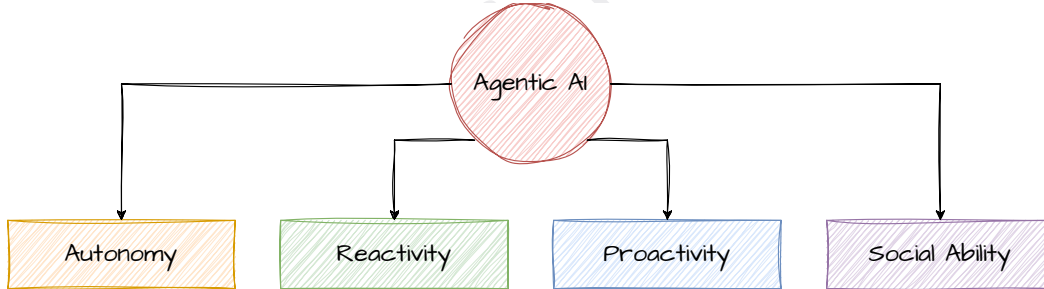


Figure 1.5: Agentic AI capabilities

- **Autonomy:** Agents independently make decisions aligned with broader business objectives without continuous human oversight.
- **Reactivity:** They rapidly detect and appropriately respond to real-time changes within their operational environment.
- **Proactivity:** They don't merely react—they proactively initiate strategies and actions aligned with pre-defined business objectives, continuously striving to achieve optimal outcomes.
- **Social Ability:** Agents can effectively communicate, collaborate, and coordinate actions with other systems, agents, and humans, working collectively toward shared objectives.

This combination of advanced capabilities positions Agentic AI as pivotal actors within the modern retail ecosystem, spanning physical retail spaces, online platforms, intricate supply chains, and diverse customer interaction points.

## 1.2 What is Agentic AI?

What	Why it matters
Agentic = autonomous, goal-directed	Goes beyond reactive generative AI
Perceive-Reason-Act-Learn loop	Mental model for readers throughout book
Combines LLMs + classic algorithms + tools	Hybrid approach yields precision and flexibility
Retail impact	Enables proactive pricing, inventory, CX decisions

At its core, **Agentic AI** refers to AI systems — often called *AI agents* — that are capable of **autonomously performing tasks** on behalf of a user or another system by dynamically designing their own workflows and using available tools (IBM Insights 2023; Russell and Norvig 2021). In other words, an Agentic AI has the *agency* to make decisions, take actions, and solve complex problems with minimal human input. Rather than being limited to pre-defined responses, these AI agents perceive their environment, reason about what they observe, and then act to achieve specified goals. They can even interact with external data sources and services beyond the data they were originally trained on (IBM Insights 2023), allowing them to adjust to real-time information and unforeseen situations.

For instance, an Agentic AI in retail might independently detect rapid sales of a product, dynamically adjust its price, reorder inventory proactively, initiate targeted marketing campaigns, and even anticipate and manage supply chain disruptions—all without requiring direct human intervention.

It’s important to note that **Agentic AI is not just generative AI with a new name**. While generative AI (like ChatGPT) focuses on producing content in response to prompts, Agentic AI is *goal-directed* and can operate autonomously over extended periods. Agentic AI systems don’t necessarily require a prompt for each action; they can chain together sequences of decisions and actions to meet a higher-level objective. In other words, generative AI is often *reactive* (it does something after you ask), whereas Agentic AI is *proactive* — it can initiate actions, adjust to changing conditions, and drive processes forward on its own. Agentic AI also tends to incorporate multiple AI techniques (LLMs, traditional algorithms, tools, etc.) to achieve precision in decision-making that pure generative models lack (IBM 2023). This means an agentic system might generate content as one step, but it will also make choices, query databases, invoke APIs, or anything else required to reach its goal. In short, **Agentic AI systems are designed for autonomous decision-making and action**, giving them a novel form of digital agency beyond the capabilities of earlier AI approaches (Wikipedia 2023).

While the concept of AI agents isn’t entirely new, classic AI literature describes an *intelligent agent* as an entity perceiving its environment and acting to achieve goals. Agentic AI expands this foundation significantly, leveraging advances like large language models (LLMs) and reinforcement learning to craft agents far more sophisticated, adaptable, and capable of managing real-world complexities (Mnih et al. 2015; Sutton and Barto 2018).

Early examples of Agentic AI include autonomous vehicles, smart assistants, and intelligent home systems. Retail, however, is uniquely positioned to benefit greatly—from AI-driven **shopping assistants** proactively assisting customers, to **automated supply chain agents** that dynamically optimize logistics, predict shortages, and streamline inventory management.

Companies like Amazon, Walmart, and Salesforce are already deploying Agentic AI beyond basic chatbots, transforming shopping experiences, dynamic pricing, inventory replenishment, and supply chain decisions. By integrating autonomy, businesses achieve faster decision-making, uninterrupted 24/7 operations, and capabilities for complex multi-step tasks impossible with traditional software or human teams alone.

The following figure depicts an architecture of an agentic retail system showing the interaction between interface, agent, intelligence, and data layers:

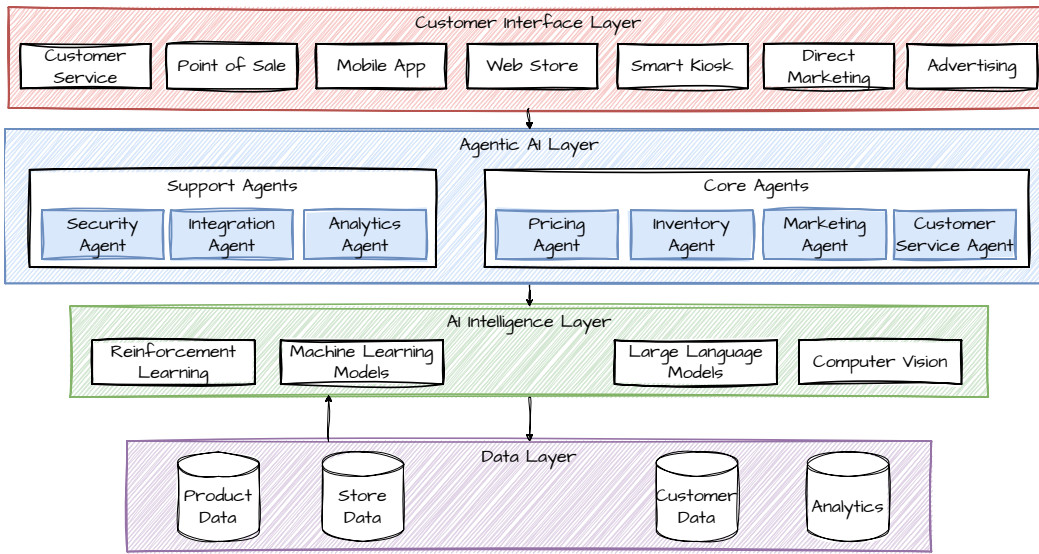


Figure 1.6: Agentic Retail System Architecture

### 1.2.1 Agentic AI vs. Traditional AI: A Paradigm Shift

Traditional AI systems typically rely on pre-programmed rules, structured datasets, and significant human intervention for decision-making. Agentic AI, in contrast, represents a new generation of AI that operates with greater *autonomy* and *adaptability*. Agentic AI learns from vast, diverse data and dynamically adjusts its behavior in real time, executing tasks **without continuous human oversight**. Instead of following static algorithms, an agentic system evolves with each interaction, improving its decision-making capabilities as it gains experience. This shift enables businesses to scale operations and respond to complexity without a proportional increase in human labor.

Agentic AI thus represents not only technological advancement but a fundamental redefinition of AI's role—from passive computational tools into active, strategic partners shaping retail's future.

### 1.2.2 How Agentic AI Works

So, how does an Agentic AI actually operate under the hood? At a high level, such an AI agent continuously goes through a **Perceive–Reason–Act–Learn** cycle (also sometimes referred to as sense–think–act or perceive–decide–act, incorporating a feedback mechanism for learning and adaptation) (Wooldridge and Jennings 1995). The 'Reason' step here encompasses planning and decision-making based on perceived information.

An Agentic AI continuously perceives data and signals feeding it into a reasoning engine. This core reasoning generates action plans executed through external APIs or tools. Outcomes then feed back as a learning signal, creating a "data flywheel" enabling continuous improvement (NVIDIA 2023). The learning phase is crucial for adaptation, allowing the agent to refine its future reasoning and actions based on past results, distinguishing it significantly from systems that only perceive and act based on fixed logic.

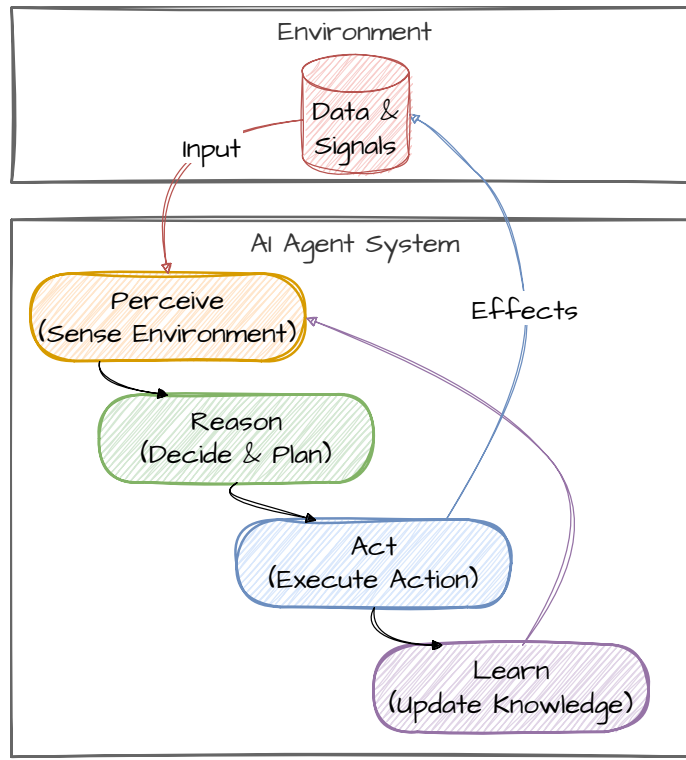


Figure 1.7: Agentic AI System Loop

An Agentic AI can be broken down into a sequence of steps or capabilities that the agent employs to function autonomously:

1. **Perception (Sensing)** – The agent gathers data from its environment and inputs. This could include real-time information from internal databases, external APIs, user interactions, sensors (if physical), etc. The goal in this step is to *perceive* the current state of the world relevant to its objectives. For example, a retail AI agent might pull the latest sales numbers, inventory levels, web analytics, or a customer's query — essentially, anything that provides context. This raw input is then processed into a form the AI can reason about (for instance, extracting key features or facts). The agent's perception component ensures it is situationally aware and working with up-to-date information.
2. **Reasoning (Planning/Deciding)** – Next, the agent analyzes the information, formulates a plan, and makes decisions. In modern Agentic AI, this often involves an LLM or other AI models acting as the *brain* of the agent. Given the goals and the perceived state, the agent generates possible solutions or actions. This may include predicting outcomes, evaluating options, and selecting the best course of action to achieve its objectives (Bratman 1987). The reasoning step is like the agent "thinking things through." For complex tasks, the agent can break the problem into sub-tasks, use specialized skills or tools (for example, calling a pricing optimization algorithm), and then assemble a solution. Thanks to advanced techniques like retrieval-augmented generation (RAG) (where an agent retrieves relevant information from external knowledge sources before generating a response or plan, enhancing accuracy and contextuality), the agent's decisions can incorporate both learned knowledge and up-to-the-minute data. The result of this phase is a decision or an action plan (e.g., *"reduce the price of item X by 10% for the next 48 hours and send a restock order for 500 units"*).

3. **Action (Execution)** – Once a plan is in place, the agent acts. It executes the chosen actions by interfacing with the necessary systems or tools. In a software context, this could mean calling APIs, updating databases, sending messages or commands – any operation that affects the environment or accomplishes a task. For a retail AI agent, actions span a wide range: adjusting a pricing database, posting a promotional campaign via a marketing API, placing an order with a supplier, or interacting with a customer through a chatbot interface. Agentic AI frameworks often integrate with external tools seamlessly, allowing the AI to, say, not only decide *what* email to send to a customer but also to go ahead and send it. It's in this stage that the AI agent tangibly impacts the business. Importantly, developers can enforce **guardrails** on actions to ensure safety and compliance. For instance, an agent might be allowed to refund a purchase up to a certain dollar amount on its own, but require human approval for anything beyond that limit. These constraints ensure the agent's autonomy remains within acceptable boundaries.
4. **Learning (Feedback Loop)** – A defining feature of Agentic AI is that it can **learn from the results of its actions** (Sutton and Barto 2018). After acting, the agent observes the new state of the environment and evaluates the outcome of its actions. Did the action succeed? Did it move closer to the goal or solve the problem? This feedback is then used to update the agent's internal knowledge or strategy for the future. Modern agent architectures implement this via a *data flywheel* or continuous improvement loop: data from interactions (e.g. how customers reacted to the price change) is fed back into the AI models, which can be retrained or fine-tuned to improve over time. In practical terms, the agent might adjust its strategy on the fly – for example, learning that certain promotions work better on weekends, or that a particular customer prefers one type of recommendation. Over many cycles, the agent becomes more effective and accurate. This adaptive ability is crucial in dynamic retail settings, where conditions and consumer behaviors are constantly changing.

These four stages — **Perceive, Reason, Act, and Learn** — form a continuous loop. The agent constantly senses the environment, thinks about what to do, does it, and then learns from what happened, then repeats. This loop enables an ongoing, autonomous operation. It's similar to how a human employee might approach a task: observe the situation, figure out a plan, execute the work, and then note what to improve next time. An Agentic AI can do this at digital speed and scale. Essentially, the agent is always asking itself: "What's going on? What should I do next? Do it. Now, how did that go and what does it mean for my next move?"

Because Agentic AI systems are quite sophisticated, they often consist of multiple sub-components or even multiple **collaborating agents**. In complex scenarios, you might have a *multi-agent system* where different agents handle different responsibilities (one focused on inventory management, another on pricing, for example) and share information with each other. They might use a shared memory or knowledge base to coordinate their efforts. However, even in these multi-agent setups, each individual agent typically follows the perceive–reason–act–learn cycle internally. The agents can negotiate or coordinate during the reasoning phase (for instance, a marketing agent may ask a supply chain agent if stock is available before launching a promo). This kind of architecture allows Agentic AI solutions to scale across various functions in an organization while maintaining autonomy and flexibility at each level (Arsanjani 2023).

It's worth noting that while current Agentic AI can adapt within predefined parameters, most do not *learn in real-time* in an unconstrained way (that could be risky). Often, the learning component involves periodic retraining or updates in controlled environments. However, as data infrastructure and AI techniques improve, we expect these agents to become increasingly self-improving in live systems. Already, the trend is toward agents that can integrate reinforcement learning or other adaptive algorithms for specialized improvements (MobiDev 2023). The end goal is an AI agent that not only automates tasks but continuously **optimizes** how it does so.

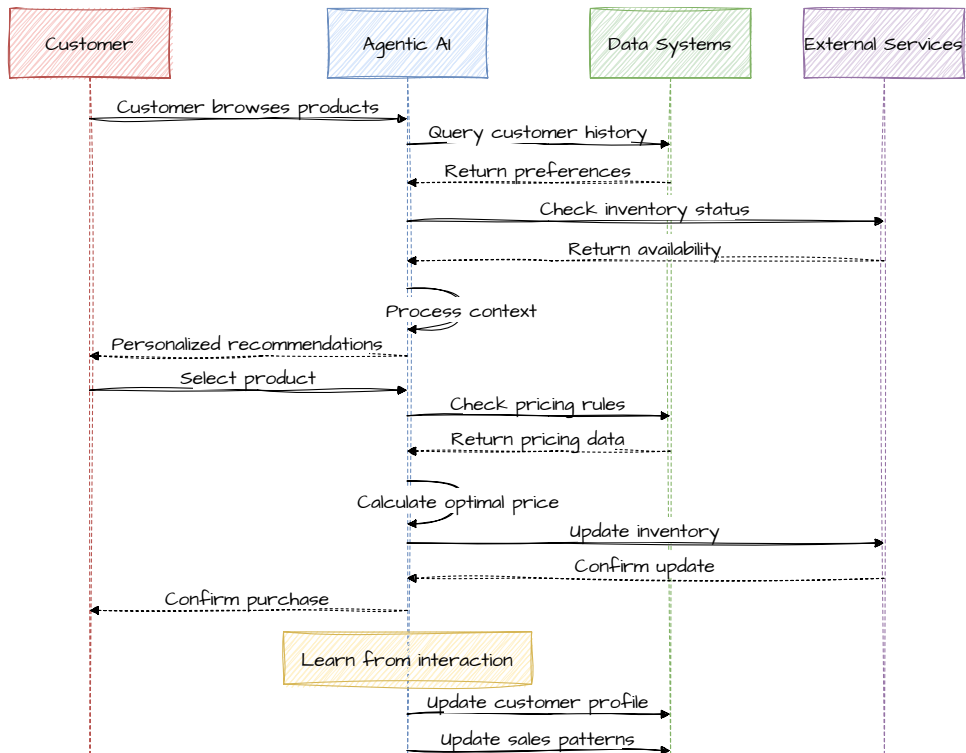


Figure 1.8: Example of interaction between customers, Agentic AI, data systems, and external services

### 1.2.3 Code Example: Implementing a Simple Agent Loop

To concretize these concepts, let's examine a simplified code example. Below is a sample python code for a very basic autonomous agent loop. This illustrative agent monitors inventory levels and decides when to reorder stock for a product. While oversimplified, it demonstrates the sense-decide-act cycle in code form.

#### **i** Code Implementation Note

The following code snippets illustrate the core concepts discussed. For the complete, executable implementation with more detailed logic and error handling, please refer to the interactive Marimo notebook for this chapter in the GitHub repository (see Preface).

```

1 # Define a simple Agent class for inventory management
2 class InventoryAgent:
3     def __init__(self, reorder_threshold, max_capacity):
4         self.reorder_threshold = reorder_threshold # When stock falls below this, agent
5         should reorder
6         self.max_capacity = max_capacity # Max storage capacity or desired stock
7         level
8         self.current_stock = 0
  
```



```

1  def perceive(self, external_data):
2      """Sense the environment: get current stock level (and any other signals)."""
3      self.current_stock = external_data.get("stock_level", self.current_stock)

```

```

1  def decide(self):
2      """Reason about whether and how much to reorder.
3
4      Implements optimal (s,S) inventory policy where:
5      - s = reorder_threshold: reorder when inventory falls below this level
6      - S = max_capacity: order up to this level when reordering
7
8      Optimality condition: s and S minimize total expected cost:
9      C(s,S) = ordering costs + holding costs + stockout costs
10     """
11     if self.current_stock < self.reorder_threshold:
12         # Plan action: calculate reorder quantity up to max capacity
13         order_quantity = self.max_capacity - self.current_stock
14         return {"action": "reorder", "amount": order_quantity}
15     else:
16         # No action needed
17         return {"action": "wait"}

```

```

1  def act(self, decision):
2      """Execute the decided action (e.g., place an order)."""
3      if decision["action"] == "reorder":
4          amount = decision["amount"]
5          print(f"Placing order for {amount} units.") # In real system, call supplier API
6          # For simulation, assume order immediately refills stock:
7          self.current_stock += amount

```

Continued in full edition