



Engineering the Agentic AI Stack



Dr. Joseph Aaron Tsapa PD

ENGINEERING THE AGENTIC AI STACK

First Edition

Author

Dr. Joseph Aaron Tsapa PD



Title of the Book: Engineering the Agentic AI Stack

First Edition – 2026

Copyright 2026 © Dr. Joseph Aaron Tsapa PD

No part of this book may be reproduced or transmitted in any form by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the copyright owners.

Disclaimer

The author is solely responsible for the contents published in this book. The publishers don't take any responsibility for the same in any manner. Errors, if any, are purely unintentional, and readers are requested to communicate such errors to the author or publishers to avoid discrepancies in the future.

E-ISBN: 978-93-7020-631-1

MRP Rs.190/-

Publisher, Printer & Distributor:

Selfypage Developers Pvt Ltd.,
Pushpagiri Complex,
Beside SBI Housing Board,
K.M. Road Chikkamagaluru, Karnataka.
Tel.: +91-8861518868
E-mail: info@iipbooks.com

IMPRINT: IIP Iterative International Publishers

For Sales Enquiries:

Contact: +91- 8861511583
E-mail: sales@iipbooks.com

PREFACE

This book is a comprehensive technical guide to agentic artificial intelligence, AI systems capable of perceiving their environment, reasoning about complex goals, planning multi-step solutions, taking autonomous action, and learning from outcomes. It is written for practitioners, researchers, policymakers, and informed technologists who need to understand not just what agentic AI can do, but also how it works, how to build it, and how to govern it responsibly.

The field of agentic AI is evolving at an extraordinary speed. Between the time this manuscript was completed and the time you read it, new models will have been released, new benchmarks achieved, new frameworks published, and new regulations proposed. This book is designed to provide durable understanding rather than ephemeral news: the formal foundations, architectural patterns, safety principles, and governance frameworks described here will remain relevant even as specific implementations evolve.

The book is organized into four parts. Chapters 1–4 establish the technical foundations: formal definitions, cognitive architectures, planning algorithms, and multi-agent coordination. Chapters 5–6 examine real-world applications in enterprise and scientific research. Chapters 7–9 address the critical challenges of safety, governance, and societal impact. Chapters 10–12 provide practical guidance for building, deploying, and understanding the long-term trajectory of agentic systems. Every chapter has been written to stand independently while building on concepts introduced earlier. Cross-references are provided throughout to help readers navigate the connections between technical architecture, safety requirements, governance frameworks, and practical implementation. Whether you are an engineer building your first agent, a researcher pushing the frontier, or a policymaker crafting governance frameworks, this book aims to equip you with the understanding needed to contribute to a future where autonomous minds serve humanity’s highest aspirations.

ACKNOWLEDGEMENT

Writing a book on Engineering the Agentic AI Stack has been both a technical and personal journey, and it would not have been possible without the support of many people.

I thank my family for their unwavering patience and encouragement through the countless late nights and weekends that went into this work. Their belief in me was the quiet engine behind every chapter.

I am grateful to my mentors and colleagues, whose conversations, critiques, and shared curiosity helped shape the ideas in these pages. Many of the frameworks and perspectives presented here were sharpened in whiteboard debates, reviews, and the honest feedback that only trusted peers can offer.

A special thanks goes to the broader open-source and research community. The builders, maintainers, and contributors behind the tools, models, and frameworks of the modern agentic stack make work like this possible, and I hope this book contributes something useful in return.

I am indebted to the reviewers and early readers who generously lent their time to examine the manuscript. Their thoughtful suggestions improved both the accuracy and clarity of these pages, and any remaining shortcomings are entirely my own.

To the editorial and publishing team, thank you for believing in this project and for the care you brought to turning a manuscript into a book.

Finally, to you, the reader: Agentic AI is evolving faster than any single book can fully capture. My hope is that the principles and practices in these pages give you a durable foundation to build upon, experiment with, and ultimately surpass. Thank you for picking this up.

Dr. Joseph Aaron Tsapa PD

TABLE OF CONTENTS

Chapter No	Chapter Name	Page No
1	The Dawn of Autonomous Intelligence ----- <i>From Mathematical Foundations to the Autonomous Age</i>	1-17
1.1	Agentic AI: A Formal Definition	1 3
1.2	Historical Trajectory: From Symbolic AI to Autonomous Agents	6
1.3	The 2024–2030 Inflection Point: Convergent Enabling Factors	9
1.4	Taxonomy of Agentic Systems: Terminology and Classification	
1.5	The Five Pillars of Agentic Intelligence	10
1.6	State of the Field: Key Players and Ecosystem Map	13
2	Architecture of an Autonomous Mind ----- <i>Cognitive Stacks, Memory, and Reasoning Frameworks</i>	18-38
2.1	The Cognitive Stack: A Six-Layer Reference Architecture	18
2.2	Foundation Models as Parametric Knowledge Engines	20
2.3	Chain-of-Thought Reasoning and Advanced Prompting Architectures	22
2.4	Memory Systems: Architecture and Engineering	24
2.5	Tool Integration: The Extended Agent Mind	26
2.6	The ReAct Framework: Formal Specification	28
2.7	Self-Reflection and Metacognitive Architecture	30
2.8	Agentic Design Patterns: A Practitioner’s Catalog	32

3	The Planning Engine	39-54

	<i>Goal Decomposition, DAGs, and Adaptive Replanning</i>	
3.1	Formal Goal Representation	39
3.2	Hierarchical Task Decomposition and DAG Planning	42
3.3	Uncertainty Management and Contingency Planning	44
3.4	Long-Horizon Planning: Maintaining Coherence over Extended Tasks	46
3.5	Adaptive Replanning: Failure Detection and Recovery	48
3.6	Planning Benchmarks and Evaluation	50
4	Multi-Agent Systems	55-69

	<i>Orchestration, Communication, and Emergent Behavior</i>	
4.1	The Case for Multi-Agent Architectures	56
4.2	Orchestration Patterns: A Formal Taxonomy	57
4.3	Agent Communication Protocols and Message Semantics	60
4.4	Emergent Behaviors: Positive and Negative	61
4.5	Consensus Mechanisms and Conflict Resolution	63
4.6	Case Study: Building a Multi-Agent Research Team	65
5	Agentic AI in the Enterprise	70-84

	<i>Vertical Applications, ROI, and Implementation</i>	
		71
5.1	The \$15 Trillion Economic Opportunity: A Quantitative Framework	73
5.2	Vertical Applications: Technical Architecture by Sector	
5.3	The Agentic Workforce: Human–AI Collaboration Models	75
5.4	Implementation Framework: Technical and Organizational Considerations	77

5.5	ROI Measurement Framework	79
5.6	Enterprise Case Study: Agentic AI in Insurance Claims Processing	81
6	Agentic AI in Science and Research	85-97

	<i>Accelerating Discovery: Drug Development, Climate Science, and the AI Scientist</i>	
6.1	The Augmented Scientific Method	86
6.2	Drug Discovery: Architecture and Empirical Results	87
6.3	Climate Science and Earth Systems	89
6.4	Materials Science and Clean Energy	91
6.5	The AI Scientist: Toward Autonomous Research	92
6.6	Practical Exercise: Designing a Literature Review Agent	94
7	Safety, Alignment, and the Control Problem	98-113

	<i>Ensuring Beneficial Autonomous Operation Across the Risk Spectrum</i>	
7.1	Why Agentic AI Safety Is Qualitatively Different	99
7.2	Alignment: Formal Failure Modes	100
7.3	Capability Risk Tiers and the Minimal Footprint Principle	102
7.4	Constitutional AI and Value Learning	103
7.5	Human Oversight: Engineering and Architecture	105
7.6	Adversarial Attacks: Prompt Injection and Defense	106
7.7	Safety Evaluation Frameworks and Benchmarks	108

8	Governance and Regulation	114-127

	<i>Legal Frameworks, Standards, and International Coordination for the Autonomous Age</i>	
8.1	The Governance Gap: Structural Analysis	115
8.2	EU AI Act: Technical Analysis	
8.3	Global Regulatory Landscape: Comparative Analysis	116
8.4	Standards Ecosystem: ISO, NIST, and IEEE	118
8.5	Liability Frameworks: The Distributed Causation Problem	120
8.6	Compliance Playbook: A Practitioner’s Guide	121
9	Agentic AI and Human Society	128-138

	<i>Cultural, Ethical, and Philosophical Dimensions</i>	
9.1	Agency: Philosophical and Computational Perspectives	129
9.2	Employment: Quantitative Impact Analysis	130
9.3	Privacy Architecture in the Agentic Era	132
9.4	Authenticity and Creative Identity	133
9.5	Digital Equity and the Access Divide	135
10	Building Agentic AI Systems	139-151

	<i>A Practical Technical Guide for Practitioners</i>	
10.1	System Design Methodology	140
10.2	Foundation Model Selection for Agentic Tasks Memory Architecture Engineering	142
10.3	Evaluation Framework for Agentic Systems	143
10.4	Production Operations: Monitoring,	144
10.5	Observability, and Iteration	146
10.6	Framework Selection Guide	147

11	The Human-Agent Partnership	152-163

	<i>Collaboration Models, Trust Calibration, and Designing for Flourishing</i>	
11.1	Reframing Human–AI Collaboration	153
11.2	Cognitive Offloading and Extended Intelligence	154
11.3	Trust Calibration: Developing Accurate Agent Models	156
11.4	The Emotional Dimension and Ethical Design	158
11.5	Designing for Human Flourishing	159
12	The Civilizational Horizon	164-173

	<i>Long-Term Trajectories, Geopolitics, and the Ultimate Question</i>	
12.1	Long-Term Scenarios: A Decision-Analytic Framework	165
12.2	The Capabilities Roadmap: 2025–2035	167
12.3	Toward Effective Global AI Governance	168
12.4	The Ultimate Question: Partnership or Succession?	169
12.5	A Call to Intentional Action	171
	Appendices	

	Appendix A: Glossary of Technical Terms	174-175
	Appendix B: Annotated Bibliography	176-177

ABOUT THE AUTHOR



Dr. Joseph Aaron Tsapa PD is a pioneering leader in Artificial Intelligence, Agentic AI Systems, Data Science, and Machine Learning, with extensive hands-on experience in transforming industries through autonomous technology solutions. With a proven track record of innovation in AI architecture and multi-agent orchestration, Joseph has developed robust agentic frameworks and introduced advanced technologies, including autonomous planning engines and cognitive AI stacks, across sectors such as banking, finance, healthcare, and utilities.

Joseph's expertise spans foundation models, reasoning architectures, tool integration protocols, and a wide array of AI frameworks and programming languages. His forward-thinking approach to Agentic AI leverages cutting-edge tools such as ReAct loops, retrieval-augmented generation, and multi-agent orchestration to automate and streamline complex processes, delivering significant cost savings and dramatic improvements in operational efficiency. Notably, his groundbreaking work on integrating Large Language Models (LLMs) into autonomous decision-making and agentic workflows has set new standards for AI capability and reliability.

Joseph is an accomplished author and researcher with more than 25 internationally published papers exploring the transformative role of artificial intelligence across banking, healthcare, and utilities. He has authored several influential books, including *Generative AI: Concepts and Applications*, *Practical Machine Learning: Real-World Applications and Techniques*, *Data Science & Machine Learning: The Modern Practitioner's Guide*, and *Strategic Leadership Velocity*.

Throughout his career, Joseph has led major AI transformation initiatives, including the deployment of multi-agent systems that have significantly reduced processing times and increased productivity. His leadership in AI safety, alignment, and governance has redefined compliance practices within regulated industries, bridging the gap between autonomous capability and responsible deployment.

Joseph holds a Professional Doctorate (PD) in Computer Science from European International University – Paris, and a Doctor of Advanced Studies (D.A.S., Honorary Dr. h.c.) in Computer Science from Azteca University. He also earned a Master of Science in Software Engineering from the Birla Institute of Technology and Science (BITS), Pilani. An active editorial board reviewer for several international journals, Joseph remains committed to advancing the field through continuous learning and knowledge sharing.

Recognized with multiple awards for his contributions, Joseph continues to set new benchmarks for Agentic AI integration and autonomous systems across data-driven industries. His visionary leadership and proven track record ensure that he will remain at the forefront of Agentic AI innovation for years to come.



SelfyPage Developers Pvt Ltd

ISBN : 978-93-7020-209-2



MRP: Rs.1000/-