Keerthi Amistapuram

From Data Pipelines to Decision Autonomy

Deep Learning and Agentic Artificial Intelligence Architectures for Intelligent Insurance Platforms

DeepScience

From Data Pipelines to Decision Autonomy: Deep Learning and Agentic Artificial Intelligence Architectures for Intelligent Insurance Platforms

Keerthi Amistapuram

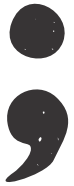

DeepScience

Published, marketed, and distributed by:

Deep Science Publishing, 2026
USA | UK | India | Turkey
Reg. No. MH-33-0658412
www.deepscienceresearch.com
editor@deepscienceresearch.com
WhatsApp: +91 7977171947

ISBN: 978-93-7185-212-8

E-ISBN: 978-93-7185-416-0

https://doi.org/10.70593/978-93-7185-416-0

Citation: Amistapuram, K. (2026). *From Data Pipelines to Decision Autonomy: Deep Learning and Agentic AI Architectures for Intelligent Insurance Platforms*. Deep Science Publishing. https://doi.org/10.70593/978-93-7185-416-0

Preface

The insurance industry stands at a pivotal crossroads. Traditional actuarial models and rule-based systems, while foundational, are increasingly insufficient to navigate the complexities of modern risk assessment, fraud detection, and customer personalization. The convergence of deep learning, cloud-native architectures, and agentic AI systems has opened unprecedented opportunities to transform insurance from a reactive, process-heavy industry into a proactive, intelligence-driven ecosystem. This book emerged from years of hands-on experience designing and deploying machine learning systems across healthcare, financial services, and retail sectors, each demanding precision, scalability, and ethical responsibility. What became evident was that the future of insurance lies not merely in automating existing workflows, but in building platforms capable of autonomous decision-making, continuous learning, and adaptive risk management.

From Data Pipelines to Decision Autonomy bridges the gap between theoretical AI concepts and practical implementation. We explore how raw, distributed data transforms into actionable intelligence through robust pipelines, how deep learning models uncover hidden patterns in claims and underwriting, and how agentic AI architectures enable systems to reason, plan, and act with minimal human intervention. Each chapter is structured to serve both technical practitioners and strategic decision-makers. You'll find architectural blueprints for building scalable data platforms on GCP and AWS, MLOps frameworks for deploying models in production, and real-world case studies demonstrating measurable business impact. We address the critical challenges of model explainability, regulatory compliance, and ethical AI deployment considerations that cannot be afterthoughts in an industry built on trust.

Whether you're a data engineer architecting your first insurance data lake, a machine learning engineer optimizing fraud detection models, or an executive navigating digital transformation, this book provides the technical depth and strategic vision needed to build truly intelligent insurance platforms.

The journey from data to autonomous decisions begins here.

Keerthi Amistapuram

Contents