

AI & ML at QCon London 2026

Sessions

Modernising Retail at Scale: Architecting a Cloud to Edge Platform for a Global Enterprise

Monday Mar 16

Modern retail depends on a technology platform that operates consistently across digital and physical channels - from customer experiences to colleague tools and the supply chain that supports them.

Jayesh Bhayani

Sr Director Technology @Tesco PLC, Board Member at Tesco Technology and Services Europe, Board Audit Committee Member of Tesco Mobile

Behind Booking.com's AI Evolution: The Unpolished Story

Monday Mar 16

It's easy to look at a mature AI platform and imagine a grand blueprint. Ours began with none. What started as a few data scientists hacking on Perl scripts and Mysql queries has grown into an AI platform that impacts millions of travel decisions every day.

Jabez Eliezer Manuel

Principal Engineer @Booking.com - Building Next-Gen AI Platform

Rethinking Your Engineering Hiring Process & Signals for the AI Era

Tuesday Mar 17

AI has distorted the signals we rely on to hire engineers. CVs are increasingly tailored, screening can be rehearsed, tech tests can look "perfect," and even system design and behavioural answers can be polished in ways that don't reflect real on-the-job judgement.

Reece Nunn

Software Engineering Manager @BBC

Blurring the Lines: Engineering & Data Teams in the Age of AI

Tuesday Mar 17

Every senior engineer knows the feeling: a model makes a bad decision, a customer complains, and suddenly you're debugging a system that spans three teams, two pipelines, and a machine learning model nobody fully owns. Where do you even start?

Lada Indra

Head of Data Platform @Pleo, Previously Head of Data @Legend and Director API Platform BI & Data @Vonage

From Copilots to Orchestrators: A 12 Week Playbook for Training Engineering Teams Using AI

Tuesday Mar 17

Most engineering teams are stuck treating AI as autocomplete. Engineers have GitHub Copilot installed (or Claude or Cursor or whatever), they're generating snippets faster, but leaders can't connect usage to business outcomes—and developers are shipping code they don't fully understand.

Krys Flores

Staff Software Engineer @Crunchyroll, Previously @Carta, @Lob, @Simple Habit, and @Nordstromrack.com|HauteLook

Team Topologies as the 'Infrastructure for Agency' with AI

Tuesday Mar 17

The book *Team Topologies Second Edition* (2025) demonstrates convincingly that organizing business and technology for fast flow of value via empowered teams produces **outsized results for enterprises worldwide**.

Matthew Skelton

CEO & Principal @Conflux, Co-Author of "Team Topologies", Leader in Modern Organizational Dynamics for Fast Flow

Teaching Engineers, Trusting AI: How Education Enabled Autonomous Code Review

Wednesday Mar 18

At Duolingo, we realized that successful AI adoption would require deliberate learning — not just access to tools. Over the past year, we scaled AI usage across 300+ engineers through intentional dogfooding programs, live training, office hours, and AI observability dashboards.

Sarah Deitke

Software Engineer @Duolingo

The Reinvention of the Dev Team

Wednesday Mar 18

I don't need to tell you that AI has changed software development forever. You know this. Whether you're positive, negative or indifferent to this change, you can't deny that the past 2 years have radically changed the role of the software developer.

Hannah Foxwell

Independent Consultant at the Intersection of Platform Engineering, Security, and AI, Founder of AI For the Rest of Us

Ethical AI Is an Engineering Problem

Wednesday Mar 18

The systems we design encode the values we choose

Clara Higuera Cabañes

Responsible AI Program Lead @BBVA - Driving Innovation & Social Impact in AI, PhD Thesis Advisor, Member of OdiselA, Previously Data Scientist @BBC

AI is an Amplifier: Scale High Performance, Not Dysfunction

Wednesday Mar 18

AI adoption in software development is nearly universal, yet the outcomes for teams are highly variable. Why do some organizations see massive productivity gains while others see their delivery stability crash? DORA's latest research provides a key insight: AI acts as an amplifier.

Nathen Harvey

Lead of DORA and Product Manager @Google Cloud

The Realities of Building an AI Native Fintech Startup

Monday Mar 16

For fifty years, \$400 trillion in professionally managed assets sat untouched by external technology. The most analytically demanding investment workflows required a combination of PhD-level quant finance and advanced engineering that no outside vendor could sustain commercially.

David Lin

Founder and CEO @Linvest21, Previously CTO @JPMorgan

Beyond Benchmarks: How Evaluations Ensure Safety at Scale in LLM Applications

Wednesday Mar 18

As LLM systems move from prototypes to production, the gap between benchmark performance and real-world reliability becomes impossible to ignore. Models that score well on benchmarks can still fail unpredictably when facing the complexity, ambiguity, and edge cases of real users.

Clara Matos

Director of Applied AI @Sword Health, Focused on Building and Scaling Machine Learning Systems

Your Agent Sandbox Doesn't Know My Authz Model: A Standard-Shaped Hole

Wednesday Mar 18

Sandboxes are the first line of defence for agentic systems: restrict the bash commands, filter the URLs, lock down the filesystem. But sandboxes operate on the syntax of requests, not the semantics of your authorization model.

Paul Carleton

Member of Technical Staff @Anthropic, Core Maintainer of MCP

Building an AI Ready Global Scale Data Platform

Wednesday Mar 18

As organizations move from single-cloud setups to hybrid and multi-cloud strategies, they are under pressure to build data platforms that are both globally available and AI-ready.

George Peter Hantzaras

Engineering Director, Core Platforms @MongoDB, Open Source Ambassador, Published Author

Explicit Semantics for AI Applications: Ontologies in Practice

Wednesday Mar 18

Modern AI applications struggle not because of a lack of models, but because meaning is implicit, fragmented, and brittle. In this talk, we'll explore how making semantics explicit (using ontologies and knowledge graphs) changes how we design, build, and operate AI systems.

Jesús Barrasa

Field CTO for AI @Neo4j

Running AI at the Edge: Running Real Workloads Directly in the

Browser

Wednesday Mar 18

Running AI today often means choosing between Anthropic or OpenAI, and accepting the cost and privacy concerns that come with shipping data to third-parties. What if there was another way?

James Hall

Founder and Director @Parallax, Author of jsPDF

Architecting AI Driven Game Creation: From Research to Production Scale Systems

Wednesday Mar 18

I spent the last two years as an architect of the next mobile gaming platform at meta, with the ambition that anyone can create, share, and play AI-generated mobile games. What I learned is that the hard part is no longer the generation — it's everything that happens after.

Danielle An

Principal Engineer / GenAI Architect @Meta, Ph.D. with 15 years of professional experience in film, MR and gaming

Ontology Driven Observability: Building the E2E Knowledge Graph at Netflix Scale

Tuesday Mar 17

As Netflix scales hundreds of client platforms, microservices, and infrastructure components, correlating user experience with system performance has become a hard data problem, not just an observability one.

Prasanna Vijayanathan

Engineer @Netflix

Renzo Sanchez-Silva

Engineer @Netflix

From S3 to GPU in One Copy: Rethinking Data Loading for ML Training

Tuesday Mar 17

ML training pipelines treat data as static. Teams spend weeks preprocessing datasets into WebDataset or TFRecords, and when they want to experiment with curriculum learning or data mixing, they reprocess everything from scratch.

Onur Satici

Staff Engineer @SpiralDB & a Core Maintainer of Vortex (LF AI & Data), Previously Building Distributed Systems @Palantir

The Ladder Is Missing Rungs: Engineering Progression When AI Ate the Middle

Tuesday Mar 17

Career progression in engineering has traditionally followed a predictable path: junior tasks teach fundamentals, mid-level work builds judgment, senior roles require synthesis across systems.

Alasdair Allan

Scientist, Author, Hacker, Maker, Journalist, CTO @Negroni Venture Studios, Interim CTO @Evaro

Building an AI Gateway Without Frameworks: One Platform, Many Agents

Tuesday Mar 17

Early AI integrations often start small: wrap an inference API, add a prompt, ship a feature. At Zoox, that approach grew into Cortex, a production AI gateway supporting multiple model providers, multiple modalities, and agentic workflows with dozens of tools, serving over 100 internal clients.

Amit Navindgi

Staff Software Engineer @Zoox

Jatin Aneja

Leading Developer Experience @Zoox, Previously Director of Site Reliability Engineering @AppLovin

Rewriting All of Spotify's Code Base, All the Time

Tuesday Mar 17

We don't need LLMs to write new code. We need them to clean up the mess we already made.

In mature organizations, we have to maintain and migrate the existing codebase. Engineers are constantly balancing new feature development with endless software upkeep.

Jo Kelly-Fenton

Engineer @Spotify

Aleksandar Mitic

Software Engineer @Spotify

The Right 300 Tokens Beat 100k Noisy Ones: The Architecture of Context Engineering

Wednesday Mar 18

Your agent has 100k tokens of context. It still forgets what you told it two messages ago.

Patrick Debois

AI Product Engineer @Tessl, Co-Author of the "DevOps Handbook", Content Curator at AI Native Developer Community

Baruch Sadogursky

DevRel Team and Context Engineering Management @Tessl AI, Co-Author of #LiquidSoftware and #DevOps Tools for #Java Developers, Java Champion, Microsoft MVP

Refreshing Stale Code Intelligence

Tuesday Mar 17

Coding models are helping software developers move faster than ever, but weirdly, the models themselves

are not keeping up. They are trained on months-old snapshots of open source code. They have never seen your internal codebase, let alone the code you wrote yesterday.

Jeff Smith

CEO & Co-Founder @Neoteny AI, AI Engineer, Researcher, Author, Ex-Meta/FAIR

Beyond Context Windows: Building Cognitive Memory for AI Agents

Tuesday Mar 17

AI agents are rapidly changing how users interact with software, yet most agentic systems today operate with little to no intelligent memory, relying instead on brittle context-window heuristics or short-term state.

Karthik Ramgopal

Distinguished Engineer & Tech Lead of the Product Engineering Team @LinkedIn, 15+ Years of Experience in Full-Stack Software Development

Reliable Retrieval for Production AI Systems

Tuesday Mar 17

Search is central to many AI systems. Everyone is building RAG and agents right now, but few are building reliable retrieval systems.

Lan Chu

AI Tech Lead and Senior Data Scientist
