

Windsor Essex GDG DevFest 2025 Sessions Abstract

Luis Serrano - The Celestial Mechanics of Attention Mechanisms

The attention mechanism is the secret sauce behind the success of transformer models like ChatGPT and DeepSeek. It enables these models to dynamically focus on the most relevant parts of a text, determining which words or phrases are most important based on their contextual relationships. In this talk, we'll explore language models through a geometric lens, imagining words as celestial bodies floating in space. The attention mechanism acts as a gravitational force, pulling these words together to form "context galaxies" where meaning emerges. We'll also dive into the roles of the Key, Query, and Value matrices, which serve as the cosmic tools that help extract and organize information from the text. No advanced mathematical background is required, just a willingness to think creatively about addition, subtraction, and the occasional multiplication. Join us for a journey through the universe of transformers, where words, gravity, and context collide!

Daria Deriagina - Pixel Perfect Is Not Dead: How I Build UI That Matches the Figma Design

What if your design looked exactly the same in the browser as it did in Figma — every pixel, every spacing, every detail?

In this talk, I'll share how I work as both a designer and a front-end developer to make that happen. I recently designed and built the Google DevFest website from scratch, and I'll walk you through my real-world process — from designing a clean, scalable UI in Figma to implementing it with code.

You'll see how I approach structure, spacing, styles, and components with both design logic and code in mind — and how that helps avoid the usual

chaos when handing off files to developers (especially when that developer is me!).

I'll also share a few of my practical rules for keeping design consistent, readable, and ready for production — whether you're working solo or in a team. This talk is especially helpful for designers who want to understand what happens after Figma, and developers who are tired of messy, unrealistic mockups.

No live coding, no fluff — just honest lessons from someone who lives on both sides.

Don Ward - GenAI by the Numbers: Signal from the Noise

Generative AI has dominated headlines, but for innovators, the critical question remains: what is the reality behind the hype? This presentation moves beyond anecdotal evidence to provide a data-driven analysis of GenAI's current and future landscape.

We will anchor our discussion in Gartner's Hype Cycle to pinpoint where GenAI stands today and what to expect as it matures from inflated expectations toward true productivity. We'll then dive deep into the numbers, exploring current enterprise adoption rates, investment trends, and performance statistics that reveal how organizations are truly using—or struggling with—this technology.

We will try to distill how to 'skate to where the puck is going, vs where it has been'.

Joe Youssouf - Oden Forge: Building a Custom Agentic Copilot with MCP + Google ADK for Next-Gen Manufacturing

Oden Forge is a production-facing LLM copilot created with the Model Context Protocol (MCP) and Google's Agent Development Kit (ADK). Forge transforms complex manufacturing production data into actionable insights through natural language workflows, making advanced analytics accessible to process engineers and other stakeholders on the factory floor alike. These capabilities support tasks such as data labeling, ad-hoc exploratory data analysis, and real-time process optimization.

At Google DevFest, Joe will explain how Oden Forge was developed and demonstrate example workflows that deliver measurable value for manufacturing customers. He will also share the challenges, successes, and lessons learned from integrating AI-powered solutions like Oden Forge into the manufacturing sector and beyond.

Umair Durrani - brandthis: Develop your personal branding & portfolio websites with LLMs

For early-career data scientists, analysts, and developers, a polished online portfolio is more than a nice addition to a resume. It can be a strong advantage in attracting the attention of recruiters, collaborators, and potential clients. A personal website with a consistent look, matching presentations, and visually coherent plots can help you stand out in a competitive market. However, building this kind of professional brand from scratch can be time-consuming and challenging, especially without design experience.

brandthis is an R package designed to make that process simple. It uses AI to help you create a complete personal brand directly from your R workflow. With just a few prompts, it can generate a logo, a color palette, and a _brand.yml file. These assets can then be applied automatically to a

Quarto-based personal website, ggplot2 themes, and Shiny apps. The result is a cohesive visual identity that works across your projects, presentations, and applications.

This talk will present two perspectives. From the user perspective, we will go through how to install brandthis, generate brand assets, set up a Quarto portfolio, and customize it to highlight your projects. From the developer perspective, we will look at how the package integrates AI to generate creative assets, how the YAML structure works, and how branding definitions are applied across R and Quarto outputs.

Keval patel - Building Your Own AI Assistant: Web Automation with Gemini

Tired of repetitive online tasks? Imagine a smart AI assistant that can navigate the web, fill out forms, and even buy products from sites like Amazon, all by understanding your natural language commands. This is the world of web automation, and with this talk, you'll learn how to build it yourself.

We'll dive into the fundamentals of creating a powerful, open-source web agent from the ground up. You'll learn to combine the advanced reasoning of Gemini with robust browser automation to create a solution that goes beyond the capabilities of tools like the ChatGPT operator. We'll explore the core concepts behind the "agentic loop": how an agent breaks down a complex goal, executes a plan, observes the results, and intelligently adapts to its environment.

Through a practical, live-coded demonstration, we'll build an agent that can perform real-world tasks, from comparing product prices to scraping

dynamic data. You'll walk away with a clear understanding of the necessary tools, the key architectural patterns, and the confidence to start building your own. Join us to unlock the future of browser automation and empower your applications with a new level of autonomy.

Ahmed Abou Gharam - Design for Reliability 2.0: How AI is Transforming Design for Reliability

As products—from cars to microchips—become more complex, traditional Design for Reliability (DfR) methods are no longer enough. Reliability engineers once relied on static test plans, statistical models, and decades of accumulated experience. Today, with the rise of AI, we are entering a new era: DfR 2.0, where reliability is designed continuously, adaptively, and intelligently.

At the heart of this transformation is the ability to leverage vast amounts of historical test data to predict the future. Instead of running repetitive and costly experiments, AI can uncover hidden patterns, accelerate life testing, and forecast potential failures before they occur. Digital twins and machine learning models can simulate years of wear in minutes, enabling smarter design choices and faster product development.

This talk will explore how AI is reshaping reliability engineering, with practical applications spanning hardware and software alike—from predictive maintenance and anomaly detection to automated test planning, software fault prediction, and reliability reporting. But success requires more than just algorithms—it demands trust, transparency, and a shift in how engineers work. It's about building systems where developers and engineers become active co-creators alongside AI systems, shaping

workflows that make both hardware and software more efficient, reliable, and future-ready.

**Nilesh Patel, Naresh Jasotani, Satish Venugopal and Don Ward -
Inventing Tomorrow: A Fireside Chat on the GenAI Revolution and Why
You Should Be Excited**

GenAI is all the rage these days. Join us as we discuss GenAI solutions that excite us (and our customers) and why the GenAI future is to be embraced.

**Salma Aly - Discover, Connect, Execute: Building API-Smart Agents with
Knowledge Graphs.**

This interactive workshop will guide participants through the fundamentals and advanced techniques of automated API testing. Attendees will learn how to design, execute, and validate API tests using industry-standard tools and frameworks, with a strong emphasis on practical, hands-on experience. Topics include understanding REST and JSON payloads, structuring automated test scripts, handling authentication, verifying responses, and integrating API testing into continuous integration/continuous deployment (CI/CD) pipelines. Participants will leave with the skills to build robust, repeatable API test suites that improve software quality, accelerate release cycles, and ensure system reliability in modern, service-oriented architectures.

**Raneem Ghalion -From Emails to Insights: Automating Customer
Service with Google Cloud & AI**

We all know how slow replies and messy communication can cost

businesses their customers. In this session, I'll walk through how Google Cloud tools—Gmail API, BigQuery, Vertex AI (Gemini), and Looker Studio—can turn customer emails into insights and instant responses.

The session is designed to be practical, light, and inspiring—something attendees can take back and apply right away.

Ben Morss - MCP demystified: how your LLM can do things in the real world

Everyone's talking about Anthropic's new Model Context Protocol (MCP), a new standard that gives your favorite LLM access to software tools. With MCP, an LLM can send emails, make Jira tickets, or browse the web. MCP makes an LLM into an agent! We'll explain how this works in detail, and we'll show you how to use MCPs of your very own. Finally, we'll build a little MCP server of our very own.

Mark Johnson - From Models to Agents: Building the Future of Enterprise AI with Google Cloud

This session explores the end-to-end journey of enterprise AI on Google Cloud. First, we'll dive into Vertex AI, covering the various features that can be used to build, deploy, and manage machine learning models at scale. (15 min) Then, we'll shift our focus to the next frontier: Google Agentspace. You'll learn how this new platform is revolutionizing enterprise productivity by transforming your organization's data into a fleet of intelligent, autonomous agents that can reason, plan, and act. (15 min)

Andrea Yzeiri - Continuity by Design: Scalable ML/AI Systems for

Longevity

In today's data-driven world, building ML and AI solutions that deliver lasting value requires more than implementing the latest tools. This session shares practical strategies for designing end-to-end systems that are scalable, reliable, and adaptable. Participants will learn how to make informed design decisions that reduce technical debt, prioritize high-impact solutions, and balance innovation with operational efficiency. The talk covers how to evaluate the right algorithms and architectures, integrate data effectively, and leverage strong mathematical foundations to create solutions that endure. Through real-world examples and actionable frameworks, attendees will gain insights into transforming complex projects into sustainable, high-value systems that drive real-world impact.

Adam Castle and John Langlois - From C++ to the C-Suite: A Fireside Chat with CTO John Langlois

Join Adam Castle, Director of Venture Services and Partnerships at WEtech Alliance, for a fireside chat with local CTO John Langlois of Scelta, as we dive into his journey from developer to the c-suite.

John will share his personal path in technology, and explore how he went from writing code to leading at the executive level, and what lessons he picked up along the way. We'll explore the realities of the CTO role through his eyes: balancing strategy with execution, navigating rapid growth, and making decisions when both the technical and business stakes are high.

In this conversation, John will reflect on the professional skills that helped him move beyond a purely technical role, including leadership, communication, and building strong teams. He'll also share the habits and routines that keep him grounded under pressure, along with the technology trends he sees shaping the future of his company and the broader industry.

Whether you're a developer curious about what leadership looks like, a founder preparing to scale, or someone interested in the evolving role of technology leaders, John's perspective offers practical insights and candid advice from someone who has walked the path, with both boots on the ground.

Frank Abbruzzese - What's Your Cybersecurity Strategy?

We will look at some Cyber Threats. What does a strong Cybersecurity Strategy look like? Share some Phishing Scams to be aware and help protect you, your data and company. Tips and best practices to keep everyone safe

Glen Yu - Secure your GKE cluster using eBPF (extended Berkeley Packet Filter) tools.

Tired of sifting through endless logs after a security event? Traditional runtime security tools typically provide insights AFTER a breach has already occurred. Good security should be proactive and not reactive. This session dives into how you can take your Google Kubernetes Engine security posture to the next level using Tetragon.

We'll explore how Tetragon, an open-source security tool, leverages the power of eBPF to provide deep, real-time security observability and enforcement directly from the Linux kernel. We'll move beyond the basics of perimeter security and demonstrate how to gain unparalleled visibility into critical, in-cluster events.

Veronica Reingold - Press A for Assistance: Making Games Accessible with AI

As video games continue to evolve in complexity and scope, accessibility, approachability, and learnability are no longer niche considerations—they're essential. In this talk, we will explore the differences between accessibility, approachability, and learnability, as well as the overlap between them. We will discuss the difficulties in developing accessible, approachable, and learnable games, why the industry lacks clear benchmarks, and the best resources and solutions for game developers. We'll dive into some examples of modern games, the challenges they pose to players, and the current assistive techniques used in industry. Then, we'll see how various AI techniques (including Utility AI, Goal-Oriented Action Planning, and Large Language Models) can dynamically assist players without compromising gameplay. This presentation will not only explain the concepts of accessibility, approachability, and learnability, but encourage creative thinking about how emerging technologies can make games more inclusive, intuitive, and fun for everyone.

Steven Rice - Bridging AI and Robotics with Python-First Frameworks

For decades, robotics software has been a significant bottleneck, characterized by low-level C++, fragmented tooling, and steep learning curves reminiscent of early machine learning. Just as Python-centric ecosystems revolutionized machine learning by prioritizing accessibility and rapid iteration, a similar paradigm shift is now transforming robotics.

This talk will explore the rise of a new generation of Python-first frameworks designed to bridge the gap between AI and robotics systems. These tools are not mere wrappers around legacy codebases; they are

fundamentally Pythonic, offering high-level abstractions and modular architectures. We will discuss how features like seamless simulation-to-reality deployment, integrated data collection pipelines, and native Gym-style interfaces are making robotics development more intuitive for the modern ML practitioner. By embracing a development philosophy that favors simplicity and interoperability, this new approach lowers the barrier to entry, accelerates research, and empowers developers to build, test, and deploy complex autonomous systems faster than ever before. Join us to see how the future of robotics is becoming more Pythonic, paving the way for the next wave of innovation.

Wen Teoh - Hack It, Build It, Launch It: Entrepreneurship for Developers

As a developer, you already have an amazing skill: you can turn an idea into something real with just your laptop. But building a cool app or project isn't always the same as building something people will actually use—or even pay for. In this talk, we'll explore how developers can turn side projects into startups, the common mistakes to watch out for, and why now is the perfect time to try. You'll leave with practical tips, real examples, and maybe even the spark for your next big idea.

Godfrey Nolan -Taming Agentic AI: Securing Deployments in the Real World

Agentic AI systems—autonomous agents capable of making decisions and taking actions—are moving rapidly from prototypes into production environments. While their potential to transform industries is enormous, they also introduce new layers of cybersecurity risk that go beyond traditional AI deployments. How do we ensure that these agents remain

trustworthy, resilient, and aligned when operating in complex, real-world conditions?

This talk explores the most pressing cybersecurity risks associated with deploying Agentic AI, including data poisoning, prompt injection, supply chain vulnerabilities, and adversarial manipulation of autonomous workflows. We will also examine real-world scenarios where attackers could exploit these systems, and the consequences organizations may face if safeguards are not in place.

Most importantly, the session will highlight practical mitigations and best practices that teams can adopt today. Topics will include secure model lifecycle management, monitoring and auditing autonomous behaviors, implementing guardrails and fallback mechanisms, and designing defense-in-depth strategies tailored for Agentic AI systems.

Attendees will leave with a clear understanding of the evolving threat landscape, as well as actionable steps to reduce risks and increase confidence when bringing Agentic AI into production.

Nicholas Schmidt - Equity, Equality, and Equations: Using Philosophy and Ethics to Guide Algorithmic Design

What does it mean for an algorithm to be “fair”? If you build an algorithm that affects others in society, you alone do not get to decide...or do you?

This talk explores how concepts of equity and equality map onto the technical world of algorithmic decisioning. We discuss how you can tie these concepts to your domain expertise and your knowledge of the world to build an algorithm that not only solves your business problem but does so in the most beneficial way possible. Attendees will gain a deeper understanding of the trade-offs between equity and equality, and how to defend their metric choices — and their models — in both technical and regulatory settings.

Safiia Mohammed - LLM Supply Chain Security: Threats, Opportunities, and Business Resilience

As Large Language Models (LLMs) become intensely integrated in business operations, their supply chains have emerged as a critical security frontier. Unlike traditional software, LLMs rely on massive training datasets, third-party APIs, fine-tuning providers, open-source model hubs, and deployment infrastructures. Each stage introduces potential vulnerabilities. Malicious behaviours can bias outputs, compromised pre-trained models may contain hidden backdoors, and insecure API integrations expose organizations to data leakage or manipulation.

However, within these threats lie opportunities. Businesses that secure their LLM supply chains can differentiate themselves by offering trustworthy, compliance-ready AI services. This aligns with regulations like the EU AI Act and Canada's AIDA, which demand transparency in AI decision-making.

Resilient organizations will view LLM supply chain security as more than risk mitigation; it is a strategic enabler. By embedding security, privacy, and accountability into AI pipelines, businesses protect customer trust, reduce

liability, and unlock competitive advantage. Enterprises that can validate what their models do, and transparently explain the how, why, and origin of their intelligence, will lead the future of AI.

Kiran Puthan - The Agentic Leap: The New Era of AI Collaboration

In this talk, we'll discover what AI agents are, exploring how they are fundamentally different from traditional AI by acting as **autonomous, goal-oriented systems**. We'll also examine the future of automation and how these agents will transform the way we work by automating complex workflows and processes.

Mina Girges - AI Meets Security: Defending Against AI-Powered Threats

Artificial Intelligence is transforming how we build, scale, and secure digital systems — but it's also redefining how attackers operate. From AI-generated phishing campaigns to deepfake-driven fraud and automated exploit discovery, the threat landscape is evolving faster than ever. For developers and organizations, the challenge isn't just adopting AI — it's building with resilience, trust, and responsibility at the core.

In this session, Mina Girges will explore the intersection of AI and cybersecurity through the lens of enterprise leadership and developer responsibility. We'll look at how AI is being weaponized, what this means for application security and identity-first design, and how organizations can adapt their security posture without slowing down innovation. Using real-world examples and lessons drawn from global industries, Mina will share actionable strategies for securing applications, APIs, and software supply chains in an AI-driven world.

Attendees will walk away with a clearer understanding of:

How AI is changing the attack surface — and the mindset required to defend it.

Practical steps developers can take to integrate responsible AI security practices into their workflows.

The leadership perspective: aligning AI adoption, Zero Trust, and governance to protect both code and business outcomes.

Whether you're a developer, architect, or leader, this talk will equip you with the insights needed to secure the next generation of digital systems — and stay one step ahead in the age of AI-powered threats.

Moshood Saka & Richard Nonso - To Be or Not to Be: Navigating the Choice Between Career and Entrepreneurship

Most of us face the big question at some point: Should I get a career or start a business? We didn't so much choose as learn to live in the messy middle, holding day jobs while building a youth-focused tech initiative that's raised over \$200,000 through grants and partnerships. In this candid co-founder talk, we'll share how two complementary lanes, one product/ops, one strategy/partnerships guided real decisions on money, timelines, and scope; how we protected momentum and trust without sacrificing the friendship; and the moments we resized ambition, said no to distractions, and kept promises to the people we serve. No playbooks, just honest lessons, failures, and receipts from balancing a paycheck and a purpose in today's fast-changing world.

Faris Alami - Start Where You Are: Utilizing the Resilient Canvas

Adversity is inevitable—but failure doesn't have to be. The businesses that not only survive but thrive in uncertain times are those that know how to assess their strengths, identify vulnerabilities, and adapt with confidence.

Join us for a dynamic session designed to help you enhance your business's resilience and prepare for whatever challenges lie ahead.

Through a mix of interactive activities, expert guidance, and collaborative discussions, you'll uncover key areas where your business may be at risk and develop practical strategies to pivot, innovate, and grow — even in the face of disruption.

Leave with a comprehensive, ready-to-implement action plan that equips you to navigate uncertainty with confidence. Don't let the unknown hold you back—invest in your business's future and gain the tools, strategies, and mindset needed to weather any storm and seize new opportunities.

Are you ready to redefine resilience?

Troy Hector - Secure your code with Scratch powered by Gemini AI

In this hands-on session, we'll explore how to use Scratch powered by Gemini AI to spot and fix vulnerabilities in your code. Together, we'll build a small project with intentional security issues, then walk through setting up Scratch AI, including installing the tool and generating an API key through **Gemini API Studio**.

From there, we'll resolve **three common security issues** step by step, giving you practical experience with how AI can strengthen your development workflow. T

Whether you're new to security or looking to enhance your existing practices, this workshop will equip you with the knowledge and tools to write safer code with the help of AI.

Amrit Mehta - Connecting Technology to Customer and Business Outcomes - Lessons from Managing Innovation in Large Corporations

Successful implementation of technology requires an enterprise mindset. Leaders need to harness innovation at the intersection of customer needs, cutting edge technologies and business results. Throughout Amrit's career at large multinationals and now within consulting, he has led R&D, technology, software product, operations and P&L teams to successfully drive transformation and innovation at start up speed. The breakout session will focus on lessons learned and key takeaways for entrepreneurs and innovators.