

THE MOSAICO CHRONICLES



Davide Di Ruscio 

Professor at University of L'Aquila, Coordinator of the PhD
program in ICT at the University of L'Aquila

Amplifying human expertise: The vision of MOSAICO and agentic software engineering

Who we are & MOSAICO's mission

I am Davide Di Ruscio, Full Professor of Computer Science at the University of L'Aquila, where I also lead the PhD program in ICT. Within the framework of the MOSAICO research initiative, the team I lead is tasked with building the technical and conceptual foundation for the next generation of agentic software engineering systems. We are moving beyond the era of simple AI coding assistants, which often function as passive, reactive tools, to create standardized, rigorous frameworks where autonomous agents act as active participants in the software development lifecycle.

Our primary mission is to establish clear, empirically backed mechanisms for organizing, evaluating, and benchmarking intelligent agents powered by Large Language Models (LLMs). We are bridging the critical gap between raw AI capabilities and practical, industrial-grade software development workflows. By doing so, we ensure that agent-driven systems are not only powerful but also reliable, predictable, and strictly aligned with the complex, evolving needs of modern software engineering.

Technical deep dive

To establish this foundation, our work addresses three fundamental technical pillars: conceptual classification, quantitative evaluation, and asset management.

First, the rapid, often chaotic proliferation of AI tools has created a confusing landscape. Terms like “agent,” “assistant,” and “copilot” are frequently used interchangeably, obscuring the underlying architectural differences. To resolve this, we have developed a comprehensive taxonomy for agentic software engineering. This system categorizes agents based on their specific levels of autonomy, the domains they operate within (e.g., single-agent pipelines versus complex, multi-agent networks), and the nature of their interaction with human developers. This common language provides the necessary ontological clarity for researchers and practitioners to systematically analyze, compare, and integrate different tools.

Second, classification is a necessary first step, but it is insufficient without objective evaluation. Our team focuses on the creation of a robust benchmarking framework centered on Key Performance Indicators (KPIs). We assess agents not merely for code generation accuracy, but for systemic metrics such as taxonomy coverage, consistency across documentation, reasoning stability, context-retrieval efficiency, and human coordination overhead. Our goal is to provide measurable benchmarks that allow us to compare agentic systems against traditional software engineering approaches.

Finally, we are developing an Agent Repository. Much like traditional software component repositories, this architecture facilitates the discovery, sharing, and reuse of pre-configured autonomous agents and workflow templates. This formalizes the protocols by which agents interact with one another, enabling developers to compose complex multi-agent systems from modular, interoperable units rather than starting from scratch for every new project.

The future of software engineering

The work being conducted in our research group is fundamentally altering the role of the software engineer. We are witnessing a transition from a paradigm where engineers write code line-by-line to one where they act as architects, orchestrating networks of autonomous agents.

By providing clear taxonomies and standardized benchmarks, our research empowers future engineers to confidently select, trust, and audit these tools for specific, high-stakes tasks. Instead of dedicating cognitive resources to repetitive maintenance or boilerplate coding, developers can pivot towards higher-level concerns: system architecture, complex requirement analysis, and critical validation.

This ensures that as software engineering becomes increasingly agentic, human expertise remains at the core, guiding the process with oversight, ethics, and strategic intent. The future of software engineering is not about replacing the human; it is about amplifying human potential through the disciplined, standardized integration of agentic intelligence.



FOLLOW US!