

THE MOSAICO CHRONICLES



Léna Bamouh

PhD Student at IMT-Atlantique (France)
& Team Member at NaoMod research team



Théo Le Calvar

Assistant Professor in Software
Engineering at IMT Atlantique

Making AI agents work as a team: inside MOSAICO's Core Coordination Team

Modern software development is increasingly assisted by AI agents. Each of them can be impressive on its own, yet each is also limited: an agent that writes code does not test it, an agent that reviews a design does not deploy it, and none of them knows what the others are doing. MOSAICO starts from the observation that the most critical problems are the ones no single agent can solve alone. Making a community of AI agents work together, reliably and transparently, is the central challenge our team tackles.

Who are we?

We are the **NaoMod team**, part of the LS2N laboratory in Nantes, France. Our research focuses on generative approaches for software engineering: how software artefacts can be described, transformed, and produced automatically, and how the processes that build them can themselves be modelled and executed.

Our role in MOSAICO

Within MOSAICO, we lead Work Package (WP) 3, whose objective is to produce a framework for the coordination and collaboration of AI agents. Our work follows two complementary branches.

The first branch brings the BDI framework (Belief, Desire, Intention) to LLM-powered agents. BDI is a long-standing model in the multi-agent systems community: an agent holds beliefs about the state of the world, pursues desires, and commits to intentions in order to achieve them. Today's LLM-based agents are usually described only by a prompt and a list of tools, which says very little about what they are actually trying to do. By giving them an explicit BDI structure, we can describe their actions and intentions far more precisely, which in turn makes them easier to select, combine and supervise.

The second branch is the development of two kinds of agents: collaboration agents, which come up with work plans, and coordination agents, which enact those plans by delegating the actual tasks to other AI agents. In other words, one figures out what should be done and in which order; the other makes it happen.

This places our work at the crossroads of the whole project. We rely on the interaction protocol defined in WP1, on the agent repository built in WP2, and on the community governance and supervision mechanisms designed in WP4.

A closer look at the coordination agent

The first thing to say about the coordination agent is that it is an agent like any other agent in MOSAICO. It is registered, described and invoked in the same way as the rest of the community. What makes it special is its role: enacting a coordination pattern, or work plan, on behalf of the community.

Concretely, enacting a plan means three things. First, finding the right agents for each task, by querying the WP2 repository rather than relying on a fixed, hardcoded list. Second, passing task descriptions, messages and artifacts between the agents involved, using the WP1 interaction protocol. Third, doing all of this while respecting the consensus policies defined in WP4, so that the execution remains supervised and accountable.

We have taken a first concrete step in this direction. Our demonstrator, a metamodel generation process, originally relied on a coordinator whose logic was written directly into the code. For our upcoming platform demonstration, we replaced it with our prototype coordination agent, which reads a workflow from a file and enacts it. The behaviour of the system is no longer buried in the implementation: it is described in an artefact that can be inspected, versioned and modified.

Our next steps are the integration with the WP4 consensus agent and with the WP2 agent repository, so that agent selection and policy enforcement become fully dynamic.

In parallel, we are working on the other end of the problem: an agent capable of decomposing a task and generating the corresponding workflow. To support this work we have gathered datasets of workflows and task descriptions, which we are now using to develop and evaluate the agent.

Impact: reasoning beyond the reach of a single agent

If MOSAICO succeeds, the benefits go well beyond our own work package. Complex tasks will be decomposed automatically into manageable pieces rather than handed as a whole to a single overloaded agent. The execution of agent workflows will become predictable and traceable, which matters enormously when AI is involved in producing software that people depend on. And agents built by different teams, with different capabilities, will be able to collaborate efficiently within the same community.

That is the shift we are working towards: from isolated assistants to a coordinated community of agents.



FOLLOW US!