



Talking to your simulator: Agentic workflows for setting up and running CO₂ storage simulations

Modelling workflows for reservoir simulation typically follow an iterative loop: setup a case, run it, inspect the output, adjust parameters, and repeat. The manual work needed between each step can be a bottleneck in the process which takes time and effort. Recent advances in agentic AI make it possible to automate much of this loop, but checks are still needed to ensure simulations are physically meaningful, not just that the code runs without errors.

In this talk we present jutul-agent, which is specifically designed to work with the JutulDarcy Reservoir Simulation tool. Given a goal in natural language, jutul-agent is able to write and execute simulation code in a live Julia session, inspecting the resulting plots and diagnostics, and iteratively correcting its own mistakes before proposing the next run. In-built tools for testing and verification ensure that the simulation results are physically and numerically valid. The resulting output includes a workspace with runnable Julia scripts which can be used and re-used independently of the jutul-agent framework.

We demonstrate this workflow on a CO₂ storage case study, working with jutul-agent to generate an ensemble of reservoir models and carry out uncertainty quantification at a CO₂ injection site. Our work shows that agentic tools have the potential to aid current CO₂ modelling workflows and make simulation software more accessible to users. However, a human expert is still required to define the modelling intent, judge the applicability of the model, accept or reject code revisions, and decide if what the agent suggests is actually useful. In this way, the expert in the loop sets the goal and judges the result.