

Autonomous factory powered by AI and flow chemistry

PIPAc (Production Intelligente de Principes Actifs) will bring the first AI-powered autonomous industrial demonstrator for active pharmaceutical ingredients (APIs) manufacturing. The outcome is high-quality API production, safer by design. The innovation relies on several key ingredients: continuous flow chemistry, artificial intelligence, and additive manufacturing.


In *flow chemistry*, chemical reactions are performed in a continuous flow, rather than in batch mode. This allows for better control and optimization of the reaction conditions, leading to **higher yields** and better product **purity**. Continuous flow chemistry approach can also reduce the amount of produced waste and reduce energy use. A small amount of any intermediate reaction products guarantees a **safer process** at all times. Additionally, this approach is easier and faster to scale-up, **reducing time to market**, and is compact. Therefore, it allows for mobile production on-demand, which is highly relevant for this application.

Artificial intelligence (AI) takes the advantages of flow chemistry one step further by adapting synthesis parameters in real-time to an ever-changing environment. For example, the flow rate is adapted to any fluctuations of the reactor temperature to guarantee the **best yield**. Available 24/7, AI will **reduce operator's stress** and thus reduce the risk of making a mistake. Another feature enabled by AI will be **predictive maintenance** of the installation.

Finally, *additive manufacturing* (3D printing) is changing the way we design the process. We manufacture chemical reactors tailored specifically to the reaction. By using customized design and materials, we improve mixing properties, heat transfer, chemical resistance, and other characteristics. This enhances **reaction**

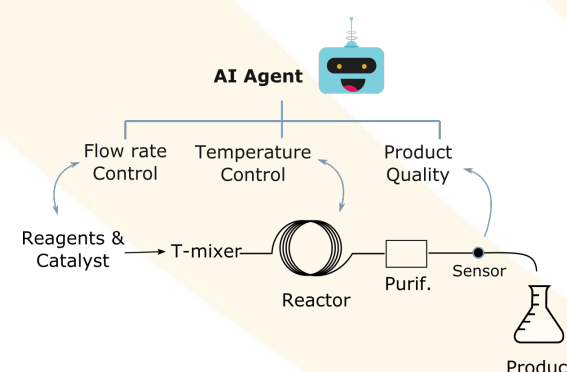
Real-World Proof of Concept

Our goal was to demonstrate for the first time the feasibility of autonomous chemical synthesis in continuous flow piloted by a deep reinforcement learning agent in a real lab. We demonstrate this on a single-step reaction of dioxolane synthesis at the temperature of 60 °C. We also imitate a non-stationarity of the environment where the agent has to adapt to decline in temperature to 40 °C or below.



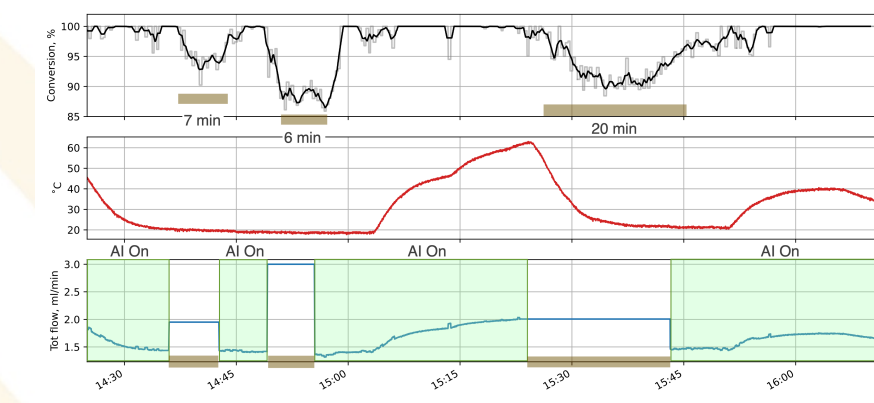
Our lab setup. Under the fume hood on the left: Continuous flow reactor, pumps, FT-IR device, thermocouples. On the right: a single-board computer hosting the AI agent.

The AI agent receives the information coming from different sensors. For example, the reactor's temperature and reaction conversion data. In its turn, the agent provides instructions to influence over controllable parameters, such as temperatures and flow rates.



For that purpose, we apply an algorithm known as *deep reinforcement learning*. The method adapts ideas from psychology and neuroscience to the context of machine learning. The AI agent only receives a reward signal compensating for its past performance. The goal is to maximize the reward. This way, the agent has to figure out optimal reaction conditions. Moreover, the AI learns to adapt

Before confronting the real-life challenge in the lab, we have trained about 1000 agents in a digital twin environment based on our past experiments. This preliminary screening allowed us to verify our key hypotheses and to optimize our AI architecture. Finally, we have selected the best 100 candidates using the best-performing neural networks. Only a few were given a chance to act under the real-world conditions.



By adapting to fluctuating environment, the AI agent achieves higher conversion compared to manual control (no AI, brown underlines).

Once running, the agent gets access to temperature readings and converted matter estimation (the reaction is analyzed in real-time by FT-IR, see [PIPAc technical note #1](#)). The agent adapts the flow rate in order to keep high product conversion. We envision that like an autopilot in airplanes, this "autopilot for chemistry" will significantly reduce operator's fatigue thus making chemical production safer and more efficient.

Conclusion and Perspectives

To our knowledge, this is **the first real-life demonstration** of chemical synthesis control in continuous flow powered by deep reinforcement learning. The method is known for previous success in playing Go and Chess and is one of components behind ChatGPT text-to-text engine training. Now, we are progressing towards adopting deep reinforcement learning to real-life challenges. Currently, we are working on chaining multiple reactions for API production.

The AI agent's ability to adjust the reaction conditions on-the-fly is particularly useful for reactions that are difficult to control or that require precise conditions to be successful. These lab-scale demonstrations will open the door towards the first industrial chemical production demonstrator autonomously piloted by AI.