

WORKSHOP

ALCF Hands-on HPC Workshop

October 7-9, 2025

Argonne National Laboratory



U.S. DEPARTMENT
of ENERGY

Argonne
NATIONAL LABORATORY



Agentic Workflows at ALCF

**Riccardo Balin, Thang Pham, Taylor Childers, Harikrishna Tummalapalli,
Aditya Tanikanti, Murat Keceli
Argonne National Laboratory**

**2025 ALCF Hands-On HPC Workshop
September 24, 2025**

What are Agentic Workflows?

- AI-powered computational pipelines that can autonomously perform complex tasks, make decisions, and adapt in real-time based on intermediate results or user input
- Agentic workflows exhibit "agency/intelligence" – can reason, plan, and execute tasks based on new results with minimal human intervention
- Key characteristics:
 - Break down complex tasks into discrete steps, each step handled by an agent
 - Autonomous and real-time decision making
 - Adaptive execution

Agentic Workflows for Science

- LLMs can provide immense value to scientific research, but struggle on complex topics (hallucinations)

Human

Provide the SMILES string corresponding to this molecule: 9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine

LLMs (with GPT-4o)

The molecule **9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine** has the following structure components:

- A **purine** core
- Substitution at position 9: a **(2,6-dichlorophenyl)methyl** group
- Substitution at position 6: an **amino group**
- That **amino group** is further substituted with a **furan-2-ylmethyl** group

Constructed SMILES:

Clc1cccc(Cl)c1CN2c3nc(nc(NCC4=CC=CO4)c3nc2)N

[...]


Incorrect

** SMILES = Simplified Molecular Input Line Entry System

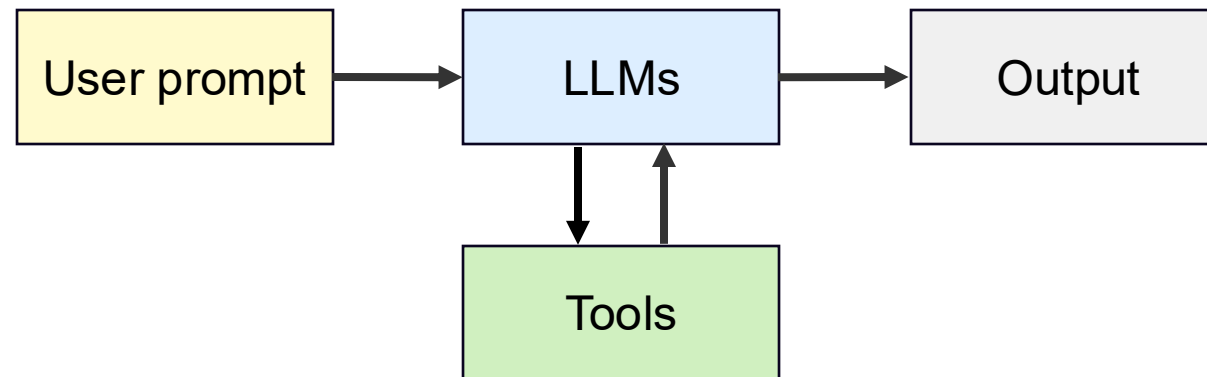
Agentic Workflows for Science

- Agentic workflows allow scientists to go beyond the “chat-bot” use of LLMs for question and answer
- They combine the power of gen-AI with computational science tools, research databases and HPC

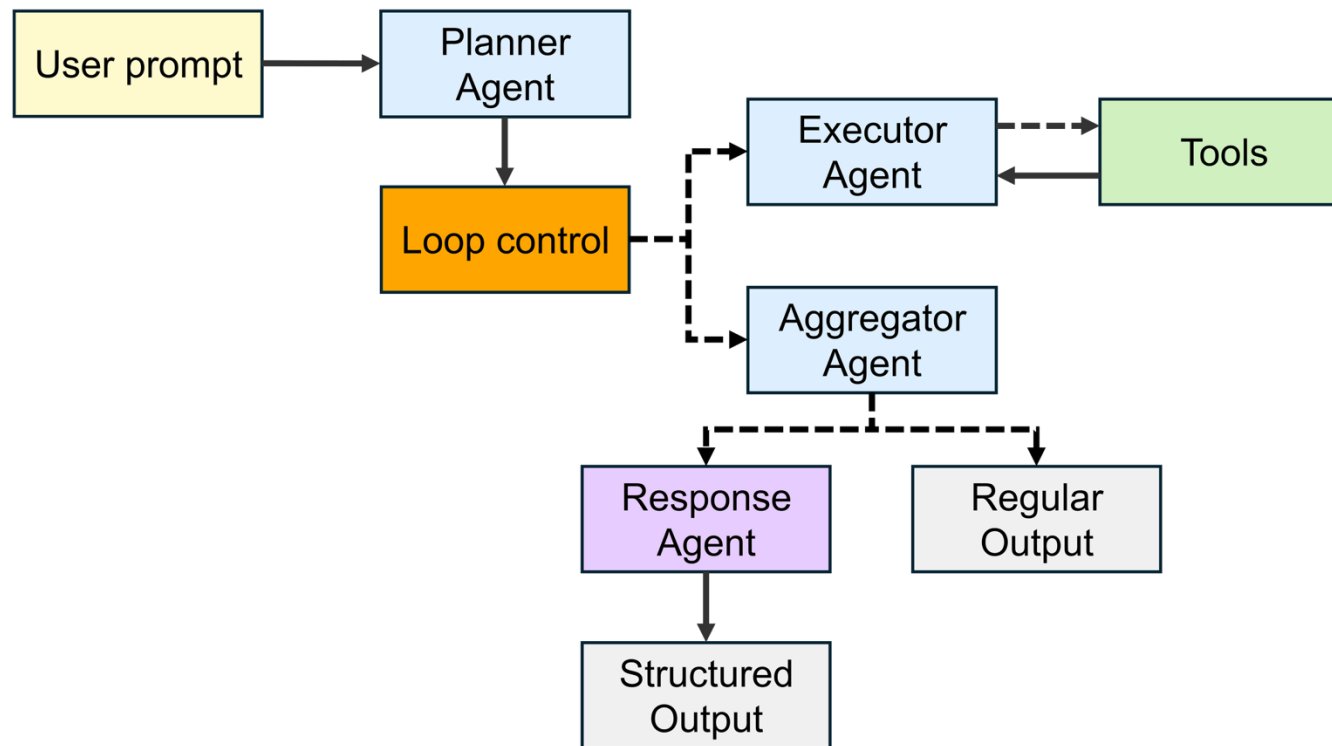
LLMs Chat-Bot



Agentic Workflow



ChemGraph: An Agentic Framework for Computational Chemistry Workflows



Software suite



ChemGraph: An Agentic Framework for Computational Chemistry Workflows. (2025)
Pham, T. D.; Tanikanti, A.; Keçeli, M. Under review. <https://arxiv.org/abs/2506.06363>

ChemGraph: Providing Accurate SMILES strings

Human

Provide the SMILES string corresponding to this molecule: 9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine

LLMs (with GPT-4o)

The molecule **9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine** has the following structure components:

- A **purine** core
- Substitution at position 9: a **(2,6-dichlorophenyl)methyl** group
- Substitution at position 6: an **amino group**
- That **amino group** is further substituted with a **furan-2-ylmethyl** group

Constructed SMILES:

Clc1cccc(Cl)c1CN2c3nc(nc(NCC4=CC=CO4)c3nc2)N

[...]

Incorrect

ChemGraph (with GPT-4o-mini)

Tool Call:

Tool name: molecule_name_to_smiles

Args: 9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine

Output:

"C1=CC(=C(C(=C1)Cl)CN2C=NC3=C(N=CN=C32)NCC4=CC=CO4)Cl"
The SMILES string corresponding to the molecule "9-[(2,6-dichlorophenyl)methyl]-N-(furan-2-ylmethyl)purin-6-amine" is

C1=CC(=C(C(=C1)Cl)CN2C=NC3=C(N=CN=C32)NCC4=CC=CO4)Cl

Correct

ChemGraph: Calculating Reaction Enthalpy

Human

You are given a chemical reaction: 1 (Methane) + 2 (Oxygen) -> 1 (Carbon dioxide) + 2 (Water). Calculate the enthalpy change for this reaction using GFN2-xTB at 400K.

ChemGraph (multi-agent, GPT-4o-mini)

Planner agent

Task 1: Calculate the enthalpy of formation of methane (CH₄) using GFN2-xTB at 400K

Task 2: Calculate the enthalpy of formation of oxygen (O₂) using GFN2-xTB at 400K

Task 3: Calculate the enthalpy of formation of carbon dioxide (CO₂) using GFN2-xTB at 400K

Task 4: Calculate the enthalpy of formation of water (H₂O) using GFN2-xTB at 400K

Executor agents

Executor 1

Executor 2

Executor 3

Executor 4

Tool calls

Tool call 1: molecule_name_to_smiles(Methane)
Tool call 2: smiles_to_atomsdata("C")
Tool call 3: run_ase(Atomsdata(Methane), driver=thermo, temperature=400, "calculator"={"calculator_type": "TBLite", "method": "GFN2-xTB"})

Tool outputs

The thermochemical calculation for methane (CH₄) using the GFN2-xTB method at 400K has been successfully completed [...]
Enthalpy of Formation: -112.2512 eV
[...]

[...]
Enthalpy of Formation: -214.93 eV
[...]


[...]
Enthalpy of Formation: -280.0425 eV
[...]

[...]
Enthalpy of Formation: -137.2903 eV
[...]

Aggregator agent


[...] Thus, the enthalpy change (ΔH) for the reaction is: $\Delta H = -12.5119$ eV

ChemGraph: Performance Across Task Complexity



Evaluation done on 13 different tasks, covering 360 instances and 200+ different molecules.

- Small language models (LLMs) are highly effective for simple, well-defined tasks.
- However, single-agent performance declines as task complexity increases.
- Multi-agent enables reaching higher accuracy even with small LLMs.



- ChemGraph UI provides a chatbot-style interface, where you can visualize results, molecular structures, vibrational modes, IR spectra, etc.
- You can run ChemGraph locally or on ALCF supercomputers.
- ChemGraph is open-source and available on GitHub



<https://github.com/argonne-lcf/ChemGraph>

Agentic Workflows at ALCF

- Users can leverage ALCF systems to accelerate scientific agentic workflows with our HPC resources
- Public GitHub with examples: <https://github.com/argonne-lcf/alcf-agentic-workflow>

Multi-system (Remote) Workflows

- Agents run and launch components on various ALCF or local systems
- LLMs served remotely (e.g., ALCF Inference Service)
- Pros:
 - Lower implementation complexity
 - Can leverage ALCF services
 - Target specific system with each computational component (efficient use of allocation)
 - Better support for long-running workflows (agents can run on local system)
- Cons:
 - Higher latency between tasks
 - Subject to scheduler when launching tasks

Single-system (Local) Workflows

- All agents and components of the workflow run on the same system, even within same batch job
- LLMs served locally on system nodes
- Pros:
 - Lower latency between tasks
 - Not subject to scheduler when running tasks
- Cons:
 - Higher implementation complexity
 - Subject to queue limitations (e.g., max. run time)
 - May consume allocation inefficiently