Do Large Language Models Speak Scientific Workflows?

Orcun Yildiz,* and Tom Peterka*
*Argonne National Laboratory, IL 60439, USA, {oyildiz,tpeterka}@anl.gov

arXiv:2412.10606v2 [cs.HC] 6 Jan 2025

***Abstract*—With the advent of large language models (LLMs), there is a growing interest in applying LLMs to scientific tasks. In this work, we conduct an experimental study to explore applicability of LLMs for configuring, annotating, translating, explaining, and generating scientific workflows. We use 5 different workflow specific experiments and evaluate several open- and closed-source language models using state-of-the-art workflow systems. Our studies reveal that LLMs often struggle with workflow related tasks due to their lack of knowledge of scientific workflows. We further observe that the performance of LLMs varies across experiments and workflow systems. Our findings can help workflow developers and users in understanding LLMs capabilities in scientific workflows, and motivate further research applying LLMs to workflows.**

***Index Terms*—Large language models, scientific workflows, code generation, workflow generation, reasoning**

I. Introduction

In recent years, we have witnessed the emergence of large language models (LLMs) increasingly being used in many fields beyond text processing. LLMs can transform scientific research by assisting researchers with complex tasks. However, current LLMs are not yet tailored for the specialized demands of scientific tasks and may struggle to perform effectively in these areas. One particularly challenging area for LLMs is scientific workflows, which are often complex, involving large number of connected tasks with wide-ranging data and computing requirements.

Over the years, the high-performance computing (HPC) community has developed workflows to support scientists in managing scientific tasks on HPC systems, meeting the demands of today's increasingly complex and large-scale computing environments. By enabling scientists to define collections of tasks along with their dependencies and data requirements, these workflows automate scheduling and data exchange of multiple workflow tasks, simplifying management of these tasks and potentially accelerating scientific discovery.

Despite their potential advantages, many scientists find workflow systems difficult to use, often opting to run tasks manually or develop their own workflow solutions instead. Even when workflow systems are adopted, scientists often lack understanding of these systems. LLMs could assist in addressing these challenges, but this requires a solid understanding of their capabilities through an extensive experimental campaign.

Several recent studies have explored the application of LLMs to specific HPC-related tasks, from code generation [1], annotation [2], answering queries [3] to scientific assistants [4], [5]. While these works provide important insights into the potential of LLMs for the HPC domain, comprehensive studies on their broader applicability within full workflow systems are still lacking.

In this work, we conduct an experimental study to explore the applicability of LLMs to HPC workflows. We evaluate five distinct workflow-specific tasks using a variety of open- and closed-source language models applied to state-of-the-art workflow systems. We hope that the insights and lessons learned from our experiments will provide a better understanding of LLM strengths and limitations in workflows, and motivate further research applying LLMs to scientific workflows.

The remainder of this paper is organized as follows. Section II presents background and related work. Section III explains our methodology in assessing the capabilities of LLMs for scientific workflows, including a description of the LLMs, workflow systems, experiments, and evaluation metrics. Section IV presents the results when applying LLMs to a series of different workflow experiments, highlighting the models' strengths and weaknesses. Section VI concludes the paper with a summary and a brief look at future work.

II. Background and Related Work

We first provide a brief background on scientific workflows and large language models. Then, we present the related work on the application of LLMs to workflows.

A. Scientific workflows

Scientific computing often relies on complex scientific workflows, which comprise a series of tasks that cooperate in terms of scheduling and communication as part of a larger scientific campaign. There are two main types of scientific workflows that we consider in this work: distributed and in situ. Distributed workflows run across several independent systems in a wide area such as grids and clouds, where the tasks exchange data through files. Examples include studies by Deelman et al. [6] and Altintas et al. [7], along with several surveys summarizing related work [8], [9]. In situ workflows, on the other hand, are designed to run within a single HPC system, with tasks executing concurrently and exchanging data through memory or the HPC system's interconnect, rather than physical storage. Representatives of such systems include ADIOS [10], Damaris [11], Decaf [12], Deisa [13], Henson [14], ParaView Catalyst [15], SENSEI [16], VisIt Libsim [17], and Wilkins [18].

B. Large language models

The introduction of transformer-based models [19] revolutionized the design of language models for understanding and generating human language. Transformer models use a self-attention mechanism to capture the dependencies in data, which allows processing data in parallel instead of sequential, improving efficiency. Large language models are trained on vast datasets from the web and contain billions of parameters. This in turn enables them to perform well in tasks such as summarization, text classification, content generation, and translation. Examples include GPT-4 [20], o1-preview, Claude 3.5 Sonnet, LLaMA [21], Bard, LaMDA [22], and PALM-2 [23].

C. Large language models for scientific workflows

Duque et al. [24] explores using LLMs to build and execute workflows. In particular, they integrate Open AI's function calling API with the Parsl workflow [25] system in order to generate workflow functions based on the user's description in natural language. Although the presented framework provides a prototype of a workflow engine that can execute a variety of tasks, it does not fully explore the challenges associated with applying LLMs to scientific workflows.

Sanger et al. [26] investigates the applicability of GPT-3.5 to three different science use cases for understanding, modifying, and extending scientific workflows. Their findings indicate that while GPT-3.5 performs well in explaining workflows, it struggles with modifying or extending workflows. Although there is some overlap with our study, their work uses GPT-3.5, which has since been followed by improved models—GPT-4o and o1-preview—that we evaluate in this study. We also provide broader range of workflow systems and scientific tasks in our study.

III. Methodology

In this section, we describe our methodology by presenting description of the LLMs, workflow systems, experiments, and evaluation metrics used in our study.

A. LLMs used for evaluation

In our experiments, we utilize several state-of-the-art large language models (LLMs), both open-source and closed-source. We selected these models because they are among the top-performing LLMs available. Below, we provide brief descriptions of each model.

- **GPT-4o and o1-preview** are closed-source language models from OpenAI. o1-preview is the latest version at the time of this writing with strong reasoning capabilities that provides step-by-step reasoning. GPT-4o [20] is the previous version before o1-preview with 1.7 trillion parameters. We use the chat interface to interact with these models.
- **Claude 3.5 Sonnet** is a closed-source language model developed by Anthropic designed for high-level problem-solving, coding, and complex language tasks. It also offers a large context window of 200,000 tokens, where GPT-4o and o1-preview models have a window size of 128,000 tokens. We use the chat interface to interact with Claude 3.5 Sonnet.
- **LLaMA-3.8B** [21] is an open-source language model from Meta with 8 billion parameters, and a context window of 128,000 tokens. LLaMA-3.8B targets enhancing natural language understanding and generation compared with former series of LLaMA models. We use the chat interface to interact with LLaMA.

B. Workflow systems

We use four different state-of-the-art workflow systems to evaluate the LLMs. Below, we provide a brief description of each system.

- **ADIOS** [10] is a middleware originally designed as a flexible I/O library for scientific codes. It has been extended to couple simulations with analysis and visualization applications through its I/O interface, facilitating seamless data exchange between different components of scientific workflows.
- **Henson** [14] is a cooperative multitasking system for in situ processing. It incorporates a built-in scripting language for workflow description and relies on shared objects and coroutines as its core abstractions. This design allows colocated tasks on the same compute node to exchange data efficiently by passing pointers directly, thereby avoiding data copies.
- **PyCOMPSs** [27] is a task-based programming model that enables the development of parallel applications and workflows. It provides task-level parallelization through the annotation of Python methods and a small API. PyCOMPSs is capable of executing on a range of distributed computing platforms, such as clusters, cloud environments, container-managed clusters, and, more recently, edge-to-cloud configurations.
- **Wilkins** [18] is an in situ workflow system that facilitates dynamic heterogeneous task specification and execution for in situ processing. It offers a data-centric API for defining the workflow graph, creating and launching tasks, and establishing communication links between the tasks.

C. Description of experiments

a) Workflow configuration: Traditionally performed manually by users prior to workflow execution, users define their workflow by indicating data and resource requirements of individual tasks using well-known languages such as YAML or Python. The workflow configuration parameter space is large and high-dimensional, making configuring workflows cumbersome, hence automating this process through LLMs can help increase productivity of workflow users.

In this experiment, we study the capability of LLMs to generate workflow configuration scripts based on natural language inputs. A sample user prompt is shown below. We evaluate the LLMs based on the correctness of the generated configuration scripts for each workflow system.

User prompt

I would like to have a 3-node workflow consisting of one producer and two consumer tasks, where producer generates grid and particles datasets, and consumer1 reads grid and consumer2 reads particles datasets. Please provide the workflow configuration file for the Wilkins workflow system.

b) Task code annotation: Today's workflow systems often require changes to the user task codes when transitioning from manually running these tasks to using a workflow system. Extensive code changes can make the learning curve steeper and create a barrier to entry into using workflows.

In this experiment, we assess the ability of LLMs to automatically annotate user task codes by asking LLMs to annotate the provided user codes for different workflow systems. Then, we evaluate the correctness of the annotated task codes. A sample user prompt is shown below.

User prompt

You are assisting in the development of a simple producer-consumer workflow using the ADIOS2 system. The producer task code is provided below. Annotate this task code in order to use it with the ADIOS2 system.

c) Task code translation: Several related works explore the applicability of LLMs for code translation between different programming languages or execution models. Similarly, our objective is assess the applicability of LLMs for task code translation between different workflow systems. For instance, users may want to compare different workflow systems for their problem, and have the task code annotated in only one system that they have been using. Ideally, users would use LLMs for manual code translation, freeing them from having to learn a new workflow system. A sample prompt is shown below, which evaluates the capability of LLMs to perform task code translation between different workflow systems.

User prompt

Task codes are provided below for the ADIOS2 workflow system for a 2-node workflow. Your task is to translate these codes to use the Henson system.

d) Workflows explanation: Today's scientific workflows often involve multiple components with diverse characteristics, making them challenging to understand when applied across various use cases. Currently, there are no established workflow standards or theoretical frameworks to guide the understanding of workflow performance, a key factor in designing robust and explainable workflows.

In this experiment, we evaluate the capability of LLMs in understanding workflows. We provide several different scientific workflow use cases and ask LLMs to explain these workflows. A sample user prompt is shown below.

User prompt

Your task is to explain the following 2-node HPC workflow. Here is the link to the producer task code: [...]. Here is the link to the consumer task code: [...] Please explain the workflow step by step.

e) Development of mini-apps for workflows: Mini-apps have been successful in other domains for characterization and understanding of applications; however, the workflow community is lacking such mini-apps. These mini-apps can emulate complex scientific workflows and be used for benchmarking different workflow systems, which can facilitate understanding of workflows.

In this experiment, we ask LLMs to develop benchmarks for workflows. In particular, we focus on the AI workflows that combine HPC and AI tasks (e.g., HPC simulation coupled with an AI analysis task). We also prompt LLMs to provide code snippets for these benchmarks that have configurable input parameters for conducting both weak and strong scaling experiments. We evaluate the LLMs by scoring the benchmarks for the soundness in the scale of 1 to 5. A sample user prompt is shown below.

User prompt

Your task is to develop benchmarks for HPC workflows. Develop 3 different benchmarks in the form of mini-apps, where the tasks are decoupled. They should represent today's scientific workflows that often combine HPC and AI tasks. Provide the code snippets for workflow tasks with benchmark parameters as inputs that are configurable to conduct both weak and strong scaling experiments.

D. Evaluation metrics

Experiments we use in our study range from code generation to understanding to reasoning. Experiments 1, 2, and 3 are code generation experiments. Experiment 4 involves understanding of the queried workflow codes. Experiment 5 involves reasoning about developing workflow benchmarks. For code generation experiments, we evaluate the LLMs based on the correctness of the generated code in the scale of 1-5. For experiments that require understanding and reasoning capabilities, we score the LLMs in terms of the soundness of their responses in the scale of 1-5.

For scoring, a full score of 5 is awarded when the response is fully functional and comprehensive—either the code executes successfully without errors or the response is sound, complete, and demonstrates a clear understanding. A score of 4 is given for responses that are mostly accurate but include minor flaws. A score of 3 reflects partial correctness, where the code or response requires noticeable adjustments or omits key aspects. A score of 2 indicates limited accuracy, with code largely incorrect or incomplete and responses showing minimal understanding. The lowest score of 1 is assigned when the response does not align with the intended workflow system or the experiment.

IV. Experiments

We evaluate the selected LLMs through a series of experiments including workflow configuration, task code annotation, task code translation, workflow interpretation, and development of mini-apps for workflows. We conduct our experiments on a workstation equipped with Apple M1 Max with 10-core CPU, 24-core GPU, 32 GB unified memory. In all of our experiments, we use the default values for model parameters.

A. Workflow configuration

In this experiment, we use the LLMs to generate workflow configuration scripts, and evaluate them based on correctness. Table I summarizes the results for each of the LLM for different workflow systems. We can see that most of the evaluated LLMs perform better for ADIOS2 and PyCOMPSs systems, while they perform poorly for Henson and Wilkins. We suspect this is mainly due to the fact that the former systems are more widely used; hence, there are more documents (scientific papers, documentation) available online for training.

Among the LLMs, we observe that o1-preview performs the best. However, we note that even o1-preview was not able to produce a correct configuration script for Wilkins and Henson systems. For example, it made up a document titled Wilkins Workflow System Documentation as a reference with the URL pointing to a personal blog website.[1] We find that LLMs struggle to infer what configuration means for scientific workflows, and often provide task code snippets instead of configuration files.

In order to highlight the importance of workflow-specific knowledge, we also report the generated configuration scripts with and without context for Wilkins using o1-preview in Table II. We provided the context by pointing o1-preview to a sample configuration file for a 2-node workflow. o1-preview generates the correct configuration file with context. Without context, one difference we can see that o1-preview uses nonexisting fields like inputs and outputs instead of inports and outports to define the data requirements. This is expected as inports/outports are specific to Wilkins, while input/output are commonly used in the literature. Also, another interesting mistake is that o1-preview calculates a dependency of consumer2 task on the consumer1 task, despite the fact that consumer2 is reading particles generated by the producer task. This demonstrates that o1-preview also lacks knowledge about in situ workflows, where tasks run concurrently, and thinks about this workflow as a pipeline where each task is executed in sequence.

B. Task code annotation

To evaluate the code annotation capabilities of LLMs, we provide a simple producer code written in C emulating an HPC simulation. We use the code available at the Henson repository[2] by replacing the Henson API with comments for LLMs to annotate with relevant workflow systems. For PyCOMPSs, we provide the equivalent of this code in Python, as PyCOMPSs is intended for Python task codes. Note that we do not involve Wilkins in this study, as Wilkins does not require modifications to task codes.

Table III shows the results for task code annotation capabilities of LLMs based on completeness of the provided response in the scale of 1-5. In overall, we see that LLMs perform better compared with the workflow configuration experiment. This can be attributed to fact that there are more code examples and documentation available online compared with the workflow configuration.

Comparing different LLMs, we see that OpenAI models outperform the others. Interestingly, GPT-4o outperforms o1 preview for ADIOS and Henson workflow systems. . For example, Henson uses $henson_yield$ for multitasking. While GPT-4o correctly adds this call at the end of the simulation code, o1-preview uses $henson_mpi_yield$, an incorrect use of Henson's API. On the other hand, o1-preview outperforms GPT-4o for PyCOMPSs, where it uses the correct call for synchronization ($compss_wait_on_file$) unlike GPT-4o ($compss_wait_on$). In particular, o1-preview is able to reason that the provided code generates a file at the end of its execution, hence, it needs to synchronize for the file name, instead of waiting on a Python object as GPT-4o attempts with $compss_wait_on$.

Even though Claude-3.5 Sonnet and LLaMA-3.8B models do not produce correct results, we are able to obtain a fully functional code after a few iterations by indicating to these LLMs the incorrect API calls. This shows the potential of these models, and highlights the need of human supervision or automated solutions as in [4] with iterative error detection and correction mechanisms.

C. Task code translation

In this experiment, we use LLMs to translate task codes between ADIOS and Henson workflow systems. We provide the annotated producer codes from the previous experiment based on the producer example from the Henson repository.

Table IV shows the results for task code translation capabilities of LLMs based on completeness of the provided response in the scale of 1-5. In overall, we observe that LLMs perform slightly worse than the task code annotation experiment. We believe this is due to involving two different workflow systems, where LLMs need to have a knowledge of both of them. Moreover, LLMs perform better when translating the producer code to ADIOS compared with translating to Henson. This is expected given the difference in the amount of online documentation for both systems.

Among the different LLMs, we again see that GPT-4o outperforms other models, while LLaMA struggles the most. We highlight the differences between the translated codes in Table V. We see that GPT-4o correctly initializes ADIOS (line 27), its IO and engine objects (lines 28-32), declares *t* and *array* variables (lines 51-52) and writes them using ADIOS's put API (lines 55-58). On the other hand, LLaMA is only able

[1] https://www.wilkins.io

[2] https://github.com/henson-insitu/henson/blob/master/examples/simple/simulation.c

	o1-preview	GPT-4o	Claude-3.5	Llama 3:8b
Workflow systems	Completeness	Completeness	Completeness	Completeness
ADIOS	4	3	3	2
Henson	1	1	1	1
PyCOMPSs	4	3	2	2
Wilkins	2	2	2	1

TABLE I
EVALUATION OF VARIOUS LLMS BASED ON THE COMPLETENESS OF THE PROVIDED WORKFLOW CONFIGURATIONS FOR EACH WORKFLOW SYSTEM.

```
tasks:
  - func: producer
    nprocs: 3
    outports:
      - filename: outfile.h5
        dsets:
          - name: /group1/grid
            file: 0
            memory: 1
          - name: /group1/particles
            file: 0
            memory: 1
  - func: consumer1
    nprocs: 5
    inports:
      - filename: outfile.h5
        dsets:
          - name: /group1/grid
            file: 0
            memory: 1
  - func: consumer2
    nprocs: 5
    inports:
      - filename: outfile.h5
        dsets:
          - name: /group1/particles
            file: 0
            memory: 1
```

```
# wilkins_workflow.yaml

tasks:
  - name: producer
    command: generate_data --output-grid grid_data --output-particles particles_data
    outputs:
      - grid_data
      - particles_data

  - name: consumer1
    command: read_grid --input-grid grid_data #nonexisting field
    inputs:
      - grid_data
    dependencies: #Wilkins automatically calculates dependencies
      - producer

  - name: consumer2
    command: read_particles --input-particles particles_data
    inputs:
      - particles_data
    dependencies:
      - consumer1 #consumer2 depends on producer as it is reading particles
```

TABLE II
GENERATED CONFIGURATION FILES WITH CONTEXT (LEFT), AND WITHOUT CONTEXT (RIGHT) USING O1-PREVIEW FOR WILKINS WORKFLOW SYSTEM. ERRORS ARE INDICATED IN RED; COMMENTS ADDED BY US TO EXPLAIN THE ERROR.

	o1-preview	GPT-4o	Claude-3.5	Llama 3:8b
Workflow systems	Completeness	Completeness	Completeness	Completeness
ADIOS	4	5	4	4
Henson	2	3	2	1
PyCOMPSs	5	4	3	2

TABLE III
EVALUATION OF VARIOUS LLMS BASED ON THE COMPLETENESS OF THE PROVIDED TASK CODE ANNOTATIONS.

to initialize the IO object correctly, and fails at other required function calls.

D. Workflows explanation

In this experiment, we provide the producer and consumer task codes that constitute a 2-node workflow, and asked LLMs to explain this workflow. We use the example task codes from each workflow system's repository. In particular, we use the heatTransfer example for ADIOS,[3] simple use case for Henson,[4] Lysozyme in water example for PyCOMPSs,[5] and ensembles use case for Wilkins.[6]

Table VI presents the results of evaluating the workflow explanation capabilities of LLMs, scored on a scale of 1-5 based on the completeness of their responses. Overall, we observe that this is the area where LLMs are most capable of, demonstrating strong performance across most experiments. Despite the good performance, hallucinations still exist, particularly when explaining workflows from less widely-used systems such as Henson and Wilkins. For instance, LLaMA inaccurately indicates that data reception at the consumer task happens via a *wilkins::receive* function call, which does not exist. Another limitation of LLMs is that they sometimes provide generic explanations based on common producer-consumer workflows rather than providing specific explanations for each queried workflow system.

Moreover, we find that o1-preview provides the most extensive explanations among the evaluated LLMs, whereas LLaMA exhibits the highest frequency of hallucinations or generic responses. An interesting observation is that both o1-preview and GPT-4o correctly infer that Wilkins uses the LowFive library for data exchange instead of files. LowFive library intercepts HDF5 API calls and redirects them to communication over the interconnect bypassing the disk, transparent to the user. Given that the provided code uses HDF5 calls and LowFive functionality is implicit in Wilkins, explaining this behavior demonstrates the advanced reasoning capabilities of these models. On the other hand, LLaMA and

[3]https://github.com/ornladios/ADIOS2/tree/master/examples/simulations/heatTransfer
[4]https://github.com/henson-insitu/henson/blob/master/examples/simple
[5]https://pycompss.readthedocs.io/en/stable/Sections/07_Sample_Applications/02_Python/05_Lysozyme_in_water.html
[6]https://github.com/orcunyildiz/wilkins/blob/master/examples/lowfive/ensembles

	o1-preview	GPT-4o	Claude-3.5	Llama 3:8b
Workflow systems	Completeness	Completeness	Completeness	Completeness
Henson to ADIOS	4	5	4	2
ADIOS to Henson	2	2	3	1

TABLE IV
EVALUATION OF VARIOUS LLMS BASED ON THE COMPLETENESS OF THE PROVIDED TASK CODE TRANSLATIONS.

Claude-3.5 Sonnet inaccurately report that the producer and consumer tasks communicate over files, misinterpreting the use of the HDF5 API in Wilkins.

E. Development of mini-apps for workflows

In this experiment, we use LLMs to develop benchmarks for workflows. We specifically asked these workflows to involve both HPC and AI tasks, and also asked to generate code snippets with configurable parameters for weak and strong scaling experiments.

Table VII presents the results of evaluating the workflow benchmark generation capabilities of LLMs, scored on a scale of 1-5 based on the soundness of their responses. Overall, we see that LLMs' capabilities are limited in this context. This experiment requires understanding both HPC workflows and benchmarking. While all LLMs suggest sound parameters for both weak and strong scaling scenarios, the provided benchmarks include application-level parameters, but lack parameters that are key to representing overall workflow performance such as the data transport. For example, o1-preview provides *grid_size* and *iterations* parameters as shown in Listing 1 (Lines 66-69), which only represent the simulation task. Ideally, such benchmarks would also include parameters for data transport in a workflow such as data transfer frequency or the amount of data exchanged between workflow tasks. Moreover, despite the fact that o1-preview indicates each benchmark involves decoupled tasks in its response (Line 8), o1-preview provides a single application where the analytics task is called within this simulation application (Line 50). These results suggest that while LLMs show promise in designing benchmarks, they lack the capability needed to address the broader challenges of developing benchmarks for scientific workflows.

F. Summary and discussion

In our experiments, the most significant limitation of LLMs we observe is their lack of understanding of scientific workflows, often leading to incorrect assumptions and hallucinations. Another limitation is the tendency of LLMs to provide generic explanations based on common producer-consumer workflow patterns rather than specific insights for queried workflow systems. We believe that addressing these limitations is vital for improving the applicability of LLMs to scientific workflows. This can be achieved by providing the contextual knowledge of scientific workflows and workflow systems to LLMs by applying several strategies, including few-shot prompting, retrieval-augmented generation (RAG), and fine-tuning. We demonstrate the effectiveness of providing this context in the workflow configuration experiment with

````
1  # Developing Benchmarks for In Situ HPC Workflows
2
3  In situ high-performance computing (HPC) workflows are essential in modern
4  scientific research, enabling simulations and data analysis to occur
5  concurrently without the overhead of writing large datasets to disk. Here,
6  we present three benchmarks in the form of mini-apps that represent
7  contemporary scientific workflows combining HPC and artificial intelligence
8  (AI) tasks. Each benchmark includes decoupled tasks and configurable
9  parameters to facilitate both weak and strong scaling experiments.
10 ---
11 ## Benchmark 1: HPC Simulation with In Situ Data Analytics
12 ### Description
13
14 This benchmark involves a parallel HPC simulation coupled with in situ
15 data analytics. The simulation generates large datasets, and the analytics
16 task processes this data in real-time to extract meaningful insights. This
17 reflects workflows in domains like climate modeling and computational fluid
18 dynamics, where simulations produce vast amounts of data that require
19 immediate analysis.
20
21 ### Workflow Tasks
22
23 - **Simulation Task**: Generates synthetic data representing the simulation output.
24 - **Analytics Task**: Performs statistical analysis on the data in situ.
25
26 ### Code Snippets
27
28 #### Simulation Task (`simulation.py`)
29
30 ```python
31 # simulation.py
32 from mpi4py import MPI
33 import numpy as np
34
35 def simulate_data(grid_size, iterations):
36     comm = MPI.COMM_WORLD
37     rank = comm.Get_rank()
38
39     # Initialize simulation data
40     local_grid_size = grid_size // comm.Get_size()
41     data = np.zeros(local_grid_size, dtype=np.float64)
42
43     for i in range(iterations):
44         # Simulate computation
45         data += np.random.random_sample(local_grid_size)
46
47         # Call analytics task
48         analytics(data, i, rank)
49
50 def analytics(data, iteration, rank):
51     # Simple statistical analysis
52     mean = np.mean(data)
53     std_dev = np.std(data)
54     print(f"Rank {rank}, Iteration {iteration},
55     Mean: {mean:.4f}, Std Dev: {std_dev:.4f}")
56
57 if __name__ == "__main__":
58     import sys
59     grid_size = int(sys.argv[1]) if len(sys.argv) &gt; 1 else 1_000_000
60     iterations = int(sys.argv[2]) if len(sys.argv) &gt; 2 else 10
61     simulate_data(grid_size, iterations)
62 ```
63
64 ### Configurable Parameters
65
66 - `grid_size`: Total size of the simulation grid (e.g., 1,000,000 points).
67 - `iterations`: Number of simulation iterations (e.g., 10, 100).
68
69 ### Scaling Experiments
70
71 - **Weak Scaling**: Increase `grid_size` proportionally
72 with the number of processes to keep workload per process constant.
73 - **Strong Scaling**: Keep `grid_size` constant while increasing
74 the number of processes to reduce computation time.
75
76 ---
````

Listing 1. One of the mini-apps generated by o1-preview. We see that this mini-app lacks the parameters that are key to representing overall workflow performance such as the data transport. o1-preview also provides a single application rather than a workflow with decoupled tasks. This highlights o1-preview's lack of understanding of HPC workflows that involve multiple interconnected tasks.

```cpp
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <time.h>
#include <mpi.h>
#include <adios2.h>

// simulation [NUM [ITER [SLEEP]]]
int main(int argc, char** argv) {
    MPI_Init(&argc, &argv);
    MPI_Comm world = MPI_COMM_WORLD;
    int rank, size;
    MPI_Comm_rank(world, &rank);
    MPI_Comm_size(world, &size);

    size_t n = 50;
    if (argc > 1) n = atoi(argv[1]);
    if (rank == 0) printf("Using %zu random numbers\n", n);

    int iterations = 3;
    if (argc > 2) iterations = atoi(argv[2]);
    int sleep_interval = 0;
    if (argc > 3) sleep_interval = atoi(argv[3]);

    srand(time(NULL) + rank);

    adios2::ADIOS adios(world);
    adios2::IO io = adios.DeclareIO("SimulationOutput");
    io.SetEngine("BPFile");
    io.AddTransport("File", {{"Library", "POSIX"}});

    adios2::Engine writer = io.Open("simulation.bp", adios2::Mode::Write);

    for (int t = 0; t < iterations; ++t) {
        if (sleep_interval) sleep(sleep_interval);

        float* array = (float*)malloc(n * sizeof(float));
        size_t i;
        for (i = 0; i < n; ++i) array[i] = (float)rand() / (float)RAND_MAX;

        float sum = 0;
        for (i = 0; i < n; ++i) sum += array[i];

        printf("[%d]: Simulation [t=%d]: sum = %f\n", rank, t, sum);

        float total_sum;
        MPI_Reduce(&sum, &total_sum, 1, MPI_FLOAT, MPI_SUM, 0, world);
        if (rank == 0) printf("[%d]: Simulation [t=%d]: total_sum = %f\n",
rank, t, total_sum);

        // Define ADIOS variables
        adios2::Variable<int> varT = io.DefineVariable<int>("t");
        adios2::Variable<float> varArray = io.DefineVariable<float>("data",
{static_cast<size_t>(size), n}, {static_cast<size_t>(rank), 0}, {1, n});

        // Write variables
        writer.BeginStep();
        writer.Put(varT, t);
        writer.Put(varArray, array);
        writer.EndStep();

        free(array);
    }

    writer.Close();
    MPI_Finalize();
    return 0;
}
```

```cpp
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <time.h>
#include <mpi.h>
#include <adios2.h>

// simulation [NUM [ITER [SLEEP]]]
int main(int argc, char** argv)
{
    MPI_Init(&argc, &argv);
    MPI_Comm world = MPI_COMM_WORLD;
    int rank, size;
    MPI_Comm_rank(world, &rank);
    MPI_Comm_size(world, &size);
    size_t n = 50;
    if (argc > 1)
        n = atoi(argv[1]);
    if (rank == 0)
        printf("Using %zu random numbers\n", n);
    int iterations = 3;
    if (argc > 2)
        iterations = atoi(argv[2]);
    int sleep_interval = 0;
    if (argc > 3)
        sleep_interval = atoi(argv[3]);
    srand(time(NULL) + rank);
    int t;
    adios2::ADIOS adios; //missing MPI_COMM_WORLD as an argument
    adios2::IO io = adios.DeclareIO("SimulationData");
    for (t = 0; t < iterations; ++t)
    {
        if (sleep_interval)
            sleep(sleep_interval);
        float* array = malloc(n * sizeof(float));
        size_t i;
        for (i = 0; i < n; ++i)
            array[i] = (float) rand() / (float) RAND_MAX;
        float sum = 0;
        for (i = 0; i < n; ++i)
            sum += array[i];
        printf("[%d]: Simulation [t=%d]: sum = %f\n", rank, t, sum);
        float total_sum;
        MPI_Reduce(&sum, &total_sum, 1, MPI_FLOAT, MPI_SUM, 0, world);
        if (rank == 0)
            printf("[%d]: Simulation [t=%d]: total_sum = %f\n", rank, t, total_sum);
        io.Write("t", t); //nonexisting API call
        io.Write("data", array, n); //nonexisting API call
        adios.PerformPuts(io); //wrong API call, engine object has the PerformPuts
        free(array);
    }
    MPI_Finalize();
    return 0;
}
```

TABLE V
TRANSLATED PRODUCER CODES WITH GPT-4O (LEFT), AND LLAMA (RIGHT) FOR ADIOS WORKFLOW SYSTEM. ERRORS ARE INDICATED IN RED; COMMENTS ADDED BY US TO EXPLAIN THE ERROR.

the o1-preview model. Moreover, for code generation tasks, we observe that with iterative user intervention—such as flagging incorrect API calls—LLMs can generate functional code despite the initial incorrect results.

Overall, we observe that the performance of LLMs varies across experiments and workflow systems. The evaluated LLMs often perform better with workflows involving well-documented systems like ADIOS2 and PyCOMPSs, likely due to the availability of relevant documentation in their training data. Our experiments show that LLMs demonstrate good performance in workflow explanation, but struggle with workflow configuration. Workflow mini-app experiment again confirms that LLMs lack understanding the complexity of HPC workflows that involve multiple interconnected tasks.

Lastly, we observe that there is no single model that is superior in all experiments. For example, OpenAI's o1-preview model demonstrates superior performance in many experiments, in particular experiments that require reasoning capabilities, but performs worse in some experiments. In contrast, GPT-4o often excels in coding-related tasks, suggesting that different models may be better suited to specific aspects of scientific workflows.

V. CONCLUSION

Scientific workflows play an important role in the scientific pipeline from hypothesis to discovery. While they offer advanced capabilities, challenges related to their usability and explainability remain. Recent advancements in LLMs present opportunities to address these issues. In this paper, we conduct an extensive experimental campaign evaluating capabilities of various LLMs across different workflow experiments and systems. Our results reveal that LLM capabilities are limited

	o1-preview	GPT-4o	Claude-3.5	Llama 3:8b
Workflow systems	Completeness	Completeness	Completeness	Completeness
ADIOS	5	5	4	4
Henson	4	2	4	2
PyCOMPSs	5	4	5	4
Wilkins	5	5	4	2

TABLE VI
EVALUATION OF VARIOUS LLMS BASED ON THE EXPLANATION OF HPC WORKFLOWS FOR EACH WORKFLOW SYSTEM.

o1-preview	GPT-4o	Claude-3.5	Llama 3:8b
3	3	2	2

TABLE VII
EVALUATION OF VARIOUS LLMS BASED ON THE SOUNDNESS OF THE PROVIDED WORKFLOW MINI-APPS.

in the context of scientific workflows, as LLMs are trained for general tasks and lack the domain knowledge required for scientific workflows. However, our results also show the promising potential of LLMs, highlighting the need for further research to address their limitations. For instance, we observe that incorporating external knowledge as context significantly improves the performance of LLMs, as we demonstrated with the o1-preview model in the workflow configuration experiment. Furthermore, we observe that iterative error correction mechanisms can be particularly effective for tasks such as scientific code generation.

ACKNOWLEDGMENT

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under contract numbers DE-AC02-06CH11357 and DE-AC02-05CH11231, program manager Margaret Lentz.

REFERENCES

[1] D. Nichols, J. H. Davis, Z. Xie, A. Rajaram, and A. Bhatele, "Can large language models write parallel code?" in *Proceedings of the 33rd International Symposium on High-Performance Parallel and Distributed Computing*, 2024, pp. 281–294.

[2] D. Nichols, A. Marathe, H. Menon, T. Gamblin, and A. Bhatele, "Hpc-coder: Modeling parallel programs using large language models," in *ISC High Performance 2024 Research Paper Proceedings (39th International Conference)*. Prometeus GmbH, 2024, pp. 1–12.

[3] L. Chen, P.-H. Lin, T. Vanderbruggen, C. Liao, M. Emani, and B. de Supinski, "Lm4hpc: Towards effective language model application in high-performance computing," in *International Workshop on OpenMP*. Springer, 2023, pp. 18–33.

[4] T. Mallick, O. Yildiz, D. Lenz, and T. Peterka, "Chatvis: Automating scientific visualization with a large language model," *arXiv preprint arXiv:2410.11863*, 2024.

[5] V. Kumar, L. Gleyzer, A. Kahana, K. Shukla, and G. E. Karniadakis, "Mycrunchgpt: A llm assisted framework for scientific machine learning," *Journal of Machine Learning for Modeling and Computing*, vol. 4, no. 4, 2023.

[6] E. Deelman, K. Vahi, M. Rynge, R. Mayani, R. F. da Silva, G. Papadimitriou, and M. Livny, "The evolution of the pegasus workflow management software," *Computing in Science & Engineering*, vol. 21, no. 4, pp. 22–36, 2019.

[7] I. Altintas, S. Purawat, D. Crawl, A. Singh, and K. Marcus, "Toward a methodology and framework for workflow-driven team science," *Computing in Science & Engineering*, vol. 21, no. 4, pp. 37–48, 2019.

[8] E. Deelman, D. Gannon, M. Shields, and I. Taylor, "Workflows and e-science: An overview of workflow system features and capabilities," *Future generation computer systems*, vol. 25, no. 5, pp. 528–540, 2009.

[9] J. Yu and R. Buyya, "A taxonomy of workflow management systems for grid computing," *Journal of grid computing*, vol. 3, pp. 171–200, 2005.

[10] D. A. Boyuka, S. Lakshminarasimham, X. Zou, Z. Gong, J. Jenkins, E. R. Schendel, N. Podhorszki, Q. Liu, S. Klasky, and N. F. Samatova, "Transparent in situ data transformations in adios," in *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, 2014, pp. 256–266.

[11] M. Dorier, G. Antoniu, F. Cappello, M. Snir, R. Sisneros, O. Yildiz, S. Ibrahim, T. Peterka, and L. Orf, "Damaris: Addressing performance variability in data management for post-petascale simulations," *ACM Transactions on Parallel Computing (TOPC)*, vol. 3, no. 3, p. 15, 2016.

[12] O. Yildiz, M. Dreher, and T. Peterka, "Decaf: Decoupled dataflows for in situ workflows," in *In Situ Visualization for Computational Science*. Springer, 2022, pp. 137–158.

[13] A. Gueroudji, J. Bigot, and B. Raffin, "Deisa: dask-enabled in situ analytics," in *2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC)*. IEEE, 2021, pp. 11–20.

[14] D. Morozov and Z. Lukic, "Master of puppets: Cooperative multitasking for in situ processing," in *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*. ACM, 2016, pp. 285–288.

[15] U. Ayachit, A. Bauer, B. Geveci, P. O'Leary, K. Moreland, N. Fabian, and J. Mauldin, "ParaView Catalyst: Enabling in situ data analysis and visualization," in *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*. ACM, 2015, pp. 25–29.

[16] U. Ayachit, B. Whitlock, M. Wolf, B. Loring, B. Geveci, D. Lonie, and E. Bethel, "The SENSEI generic in situ interface," in *Proceedings of the 2nd Workshop on In Situ Infrastructures for Enabling Extreme-scale Analysis and Visualization*. IEEE Press, 2016, pp. 40–44.

[17] T. Kuhlen, R. Pajarola, and K. Zhou, "Parallel in situ coupling of simulation with a fully featured visualization system," in *Proceedings of the 11th Eurographics Conference on Parallel Graphics and Visualization (EGPGV)*, 2011.

[18] O. Yildiz, D. Morozov, A. Nigmetov, B. Nicolae, and T. Peterka, "Wilkins: Hpc in situ workflows made easy," *arXiv preprint arXiv:2404.03591*, 2024.

[19] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.

[20] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[21] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.

[22] R. Thoppilan, D. De Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H.-T. Cheng, A. Jin, T. Bos, L. Baker, Y. Du *et al.*, "Lamda: Language models for dialog applications," *arXiv preprint arXiv:2201.08239*, 2022.

[23] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen *et al.*, "Palm 2 technical report," *arXiv preprint arXiv:2305.10403*, 2023.

[24] A. Duque, A. Syed, K. V. Day, M. J. Berry, D. S. Katz, and V. V. Kindratenko, "Leveraging large language models to build and execute computational workflows," *arXiv preprint arXiv:2312.07711*, 2023.

[25] Y. Babuji, A. Woodard, Z. Li, D. S. Katz, B. Clifford, R. Kumar, L. Lacinski, R. Chard, J. M. Wozniak, I. Foster *et al.*, "Parsl: Per-

vasive parallel programming in python," in *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*, 2019, pp. 25–36.

[26] M. Sänger, N. De Mecquenem, K. E. Lewińska, V. Bountris, F. Lehmann, U. Leser, and T. Kosch, "Large language models to the rescue: Reducing the complexity in scientific workflow development using chatgpt," *arXiv preprint arXiv:2311.01825*, 2023.

[27] E. Tejedor, Y. Becerra, G. Alomar, A. Queralt, R. M. Badia, J. Torres, T. Cortes, and J. Labarta, "PyCOMPSs: Parallel computational workflows in Python," *The International Journal of High Performance Computing Applications*, vol. 31, no. 1, pp. 66–82, 2017.