CWE-8
June 7-11, 2026
London, ON, Canada

Autonomous CFD Preparation: Agentic AI for Mesh Generation and Turbulence Representation

Fei Ding[a], Teng Wu[b]

[a]University of Utah, Salt Lake City, Utah, USA, fei.ding@utah.edu
[b]University at Buffalo, Buffalo, New York, USA, tengwu@buffalo.edu

Summary

Computational Fluid Dynamics (CFD) plays a critical role in modeling wind effects on structures and built environment. Nonetheless, the effective use of CFD, particularly with open-source software, requires substantial domain expertise in mesh generation, turbulence model selection, and boundary condition specification. These requirements limit broader adoption and often compromise simulation reliability. Recent advances in Large Language Model (LLM)–based agents offer a promising solution to automate CFD workflows through natural languages. Existing frameworks, such as MetaOpenFOAM and OpenFOAMGPT, demonstrate this potential but often struggle to achieve mesh accuracy for different CFD model fidelities and in physically sensitive regions including near-wall boundary layers. To address this gap, we propose a novel agent-based framework featuring a hierarchical Retrieval-Augmented Generation (RAG) pipeline tailored for advanced meshing tasks. The framework explicitly incorporates physical knowledge and CFD best practices to adaptively refine meshes, enabling more reliable and physically consistent simulations.

Keywords: *autonomous CFD workflows, Large Language Models (LLMs), mesh generation, Retrieval-Augmented Generation (RAG), physics-aware meshing*

1 INTRODUCTION

Setting up a single CFD simulation typically requires deep domain expertise, specialized technical skills, and substantial manual effort. Although open-source software such as OpenFOAM is a mature and widely used platform, its reliance on expert knowledge for preparing configuration files and debugging through command-line interfaces limits CFD accessibility for non-experts. Recently, rapid advances in LLMs have enabled the development of intelligent agents capable of interpreting natural-language commands and performing automated reasoning, opening new opportunities to automate CFD workflows.

The current state-of-the-art in LLM-driven CFD automation primarily leverages RAG to integrate CFD domain knowledge with general-purpose language models. RAG enables access to databases of existing OpenFOAM tutorials, providing relevant examples that guide case setup and improve modeling precision. Early pioneering systems, such as MetaOpenFOAM [1] and OpenFOAMGPT [2], demonstrated the ability to interpret natural-language instructions, coordinate multi-step tasks, and automatically generate OpenFOAM configuration files. More advanced RAG-based frameworks, including ChatCFD [3] and Foam-Agent [4], have further enhanced the accuracy and relevance of retrieved information, strengthening domain knowledge integration. In contrast to RAG-based approaches, some research has focused on embedding domain expertise directly into the model's parameters through fine-tuning, such as NL2FOAM [5]. Collectively, these systems significantly lower technical barriers, improve efficiency, and pave the way toward a new paradigm of autonomous CFD analysis.

While existing agents demonstrate strong capability in automating the overall configuration of CFD cases, a critical gap remains in the automated generation of computational meshes, which are essential for accurately simulating wind flows around civil structures. Mesh generation is not merely a geometric task but a physics-driven process that requires an understanding of flow physics, turbulence modeling, and boundary layer theory. As a result, mesh quality needs to be coupled with turbulence model fidelity and adaptively refined in physically sensitive regions, such as near-wall boundary layers where steep gradients occur. Inappropriate mesh resolution in these regions can lead to physically incorrect predictions of flow separation and aerodynamic loads on structures. This gap highlights the need for a more intelligent, physics-aware meshing approach tailored to civil engineering CFD applications.

2 METHODOLOGY

To address the challenge of automated meshing quality, we propose a novel framework centered on a specialized agent equipped with a hierarchical RAG pipeline that learns how to generate meshes by leveraging established expert-level CFD tutorials, research publications, and best-practice guidelines, while also performing physics-aware fine-tuning to adapt the mesh to flow characteristics.

2.1 Hierarchical retrieval for physics-aware meshing

The core innovation of this framework is the development of physics-aware agents that automate mesh generation while aligning with the user's specified inflow conditions, turbulence model selection, and the numerical accuracy required by the underlying fluid flow physics. The agent is built upon a multi-level retrieval mechanism that equips it with both high-level strategic guidance and detailed, context-specific information. It integrates two retrieval components: a reference retriever, which accesses a knowledge base of CFD best-practice guidelines and tutorials to provide high-level meshing strategies, and a context retriever, which extracts case-specific constraints such as geometry, inflow conditions from multi-modal user inputs.

The reference retriever searches a curated knowledge base that includes OpenFOAM tutorials, validated research cases, and CFD best-practice documents. Its role is to identify high-level meshing templates that are similar to the user's problem. For example, for a rectangular building, it may retrieve a structured mesh layout along with general guidelines for boundary layer refinement.

The context retriever focuses on fine-grained, case-specific information needed for detailed meshing decisions. For instance, it extracts the selected turbulence model, wall functions, and boundary conditions directly from the user's input. Using retrieved parameters, the agent adaptively refines the mesh to satisfy the physical requirements of wall functions in either RANS or LES turbulence models. This refinement loop continues until mesh quality metrics meet the prescribed physical and numerical constraints.

2.2 Multi-model inputs

A key feature of the agent is its ability to process multi-modal user inputs. The proposed framework accepts natural-language prompts describing geometry and inflow conditions, case studies from research papers that specify numerical setups, and existing external mesh files, such as meshes generated with Gmsh, which can serve as starting points for iterative refinement and mesh quality improvement.

3 CASE STUDY & VISUALIZATION OF RESULTS

To validate the capabilities of the proposed framework, we will conduct a demonstrative case study simulating wind flows around tall buildings with varying geometries and using different levels of CFD model fidelity. Similar to the visualizations provided in the OpenFOAM tutorial materials, the automated CFD framework will generate textual descriptions of the velocity contours and pressure distribution plots. These descriptions directly reflect mesh quality and the mesh's ability to accurately resolve the underlying flow physics, thereby providing a built-in validation of the proposed framework's effectiveness.

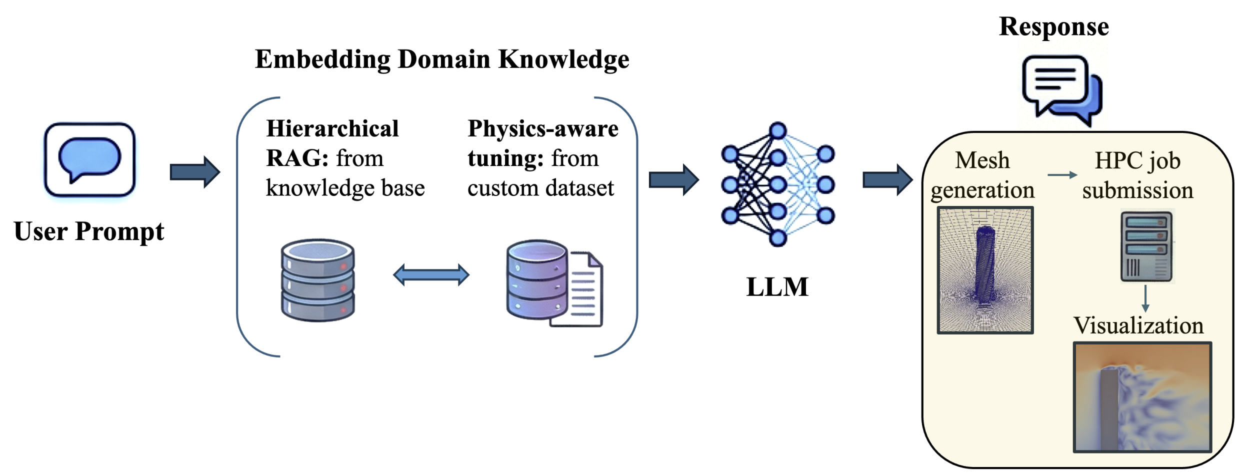

Figure 1: Workflow of the automated CFD pipeline with domain-knowledge embedded mesh generation.

4 CONCLUSIONS

This work addressed a primary bottleneck in the automated application of CFD: the lack of accuracy and physical interpretability in mesh generation produced by existing LLM-based agents. We introduce a hierarchical RAG mechanism with dedicated reference and context retrievers to overcome this limitation by iteratively updating the mesh in physics-critical regions. By systematically injecting domain knowledge from both best-practice guidelines and case-specific user input, our agent can perform intelligent, iterative refinement of the computational mesh. The overarching goal is to democratize access to CFD by enabling advanced simulation capabilities for a broader range of engineering problems that were previously intractable for non-experts.

ACKNOWLEDGEMENTS

The acknowledgement text should be styled appropriately. Acknowledge funding sources, collaborators, and any other contributions to the work.

REFERENCES

[1] Yuxuan Chen, Xu Zhu, Hua Zhou, and Zhuyin Ren. Metaopenfoam: an llm-based multi-agent framework for cfd. *arXiv preprint arXiv:2407.21320*, 2024.

[2] Sandeep Pandey, Ran Xu, Wenkang Wang, and Xu Chu. Openfoamgpt: A retrieval-augmented large language model (llm) agent for openfoam-based computational fluid dynamics. *Physics of Fluids*, 37(3), 2025.

[3] E Fan, Weizong Wang, and Tianhan Zhang. Chatcfd: an end-to-end cfd agent with domain-specific structured thinking. *arXiv preprint arXiv:2506.02019*, 2025.

[4] Ling Yue, Nithin Somasekharan, Tingwen Zhang, Yadi Cao, and Shaowu Pan. Foam-agent 2.0: An end-to-end composable multi-agent framework for automating cfd simulation in openfoam. *arXiv preprint arXiv:2509.18178*, 2025.

[5] Zhehao Dong, Zhen Lu, and Yue Yang. Fine-tuning a large language model for automating computational fluid dynamics simulations. *Theoretical and Applied Mechanics Letters*, page 100594, 2025.