

OpenFOAM Case Generation and Traversal using LLMs

Vedant Dubey

B.Tech - Electronics and Communications Technology
National Institute of Technology, Raipur

Suyash Mishra

Department of Electronics and Communications Engineering
Indian Institute of Information Technology, Manipur

Abstract

OpenFOAM is a widely used open-source toolbox for computational fluid dynamics, but its manual configuration process—requiring the precise editing of numerous complex files—creates significant barriers for users. This paper presents PyVnt, an artificial intelligence-driven toolkit developed to automate and simplify OpenFOAM case setup and management. PyVnt comprises two main components: FoamGen, a command-line utility that generates case files from natural language prompts using large language models, and a Bidirectional Converter that enables seamless transformation between OpenFOAM dictionaries and programmable tree structures. The system incorporates a retrieval-augmented generation approach, leveraging a curated knowledge base of over 800 validated case files to ensure contextually accurate outputs. A web interface built with Flask and real-time interaction via SocketIO further enhance user accessibility. Comprehensive testing demonstrates that PyVnt reduces setup time by 75–80% and supports a wide range of workflows, making advanced simulations more accessible for educational, research, and industrial users. The report details the system’s architecture, implementation, testing, and performance, concluding that PyVnt significantly streamlines CFD workflows and lowers the entry barrier for OpenFOAM users.

Acknowledgement

We would like to express our sincere gratitude to **Prof. Chandan Bose** and **Mr. Diptangshu Dey** for their guidance and support during our summer fellowship under the OpenFOAM GUI project at FOSSEE IIT Bombay. Their expertise and knowledge in the field of computational fluid dynamics and software development have been invaluable to us, and we are truly grateful for their willingness to share their time and insights with us.

We would also like to thank the staff at FOSSEE IIT Bombay for their warm welcome and support during our fellowship. It was a pleasure to work with such a talented and dedicated team.

This report would not have been possible without the support of all of the people mentioned above. We are deeply grateful for their contributions.

Table of Contents

1	Background and Context	6
1.1	Computational Fluid Dynamics Fundamentals	6
1.2	OpenFOAM Ecosystem	6
1.3	AI in CFD Automation	6
1.4	FOSSEE Initiative	6
2	System Design Principles	7
3	System Architecture	8
3.1	Data Layer	8
3.2	Processing Layer	9
3.3	AI Layer	9
3.4	API Layer	9
3.5	Presentation Layer	9
3.6	Scalability and Performance Optimization	9
3.7	Fault Tolerance	10
4	Project Management	11
4.1	Agile Methodology	11
4.2	Collaboration and Version Control	11
5	Methodology	12
5.1	Knowledge Base Construction	12
5.2	LLM Selection and Tuning	12
5.3	Testing Framework Design	13
6	Implementation Details	14
6.1	FoamGen: AI-Powered Case Generation	14
6.2	Bidirectional Converter	15
6.3	PyVnt Tree Structure	15
6.4	Web Interface	15
6.5	Knowledge Base Management	16
7	User Workflows	17
7.1	Educational Workflow	17
7.2	Research Workflow	17
7.3	Industrial Workflow	17
8	Case Studies	18
8.1	Lid-Driven Cavity	18
8.2	Multiphase Dam Break	19
8.3	Turbulent Pipe Flow	19
8.4	Heat Exchanger Simulation	20
8.5	Combustion Simulation	20

8.6 Particle Tracking	21
9 Testing and Validation	22
9.1 Validation Strategies	22
9.2 Performance Metrics	23
10 Performance Benchmarking	26
11 Limitations and Future Work	27
11.1 Current Limitations	27

List of Figures

1	System architecture of the PyVnt toolkit, illustrating component interactions and data flow.	8
2	Knowledge base structure and processing pipeline.	9
3	Comprehensive workflow efficiency comparison demonstrating 75-80% time reduction across varying complexity levels, with greatest benefits observed for advanced multiphysics simulations.	24
4	Error distribution analysis across different configuration categories, revealing parameter specification as the most challenging aspect of automated case generation. .	25

1 Background and Context

1.1 Computational Fluid Dynamics Fundamentals

Computational Fluid Dynamics (CFD) is a cornerstone of modern engineering, enabling numerical simulation of fluid flow, heat transfer, and multiphysics phenomena. CFD is governed by the Navier-Stokes equations, which describe the conservation of mass, momentum, and energy:

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} + \mathbf{f}, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2)$$

where ρ is density, \mathbf{u} is velocity, p is pressure, μ is viscosity, and \mathbf{f} represents external forces. These equations, coupled with appropriate boundary conditions, enable simulations of complex phenomena in aerospace, automotive, and environmental engineering (10).

1.2 OpenFOAM Ecosystem

OpenFOAM is an open-source C++ toolbox renowned for its flexibility in solving CFD problems, supporting solvers for laminar and turbulent flows, multiphase interactions, combustion, and more (1). Its modular structure allows customization, but configuring 10-15 dictionary files (e.g., `controlDict`, `fvSchemes`, `fvSolution`) requires precise syntax and domain expertise, leading to errors and prolonged setup times.

1.3 AI in CFD Automation

The integration of artificial intelligence, particularly Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG), offers a transformative approach to automating CFD workflows. LLMs can interpret natural language prompts, while RAG enhances accuracy by retrieving contextually relevant data (4). This project leverages these technologies to abstract OpenFOAM's complexities, making CFD accessible to a broader audience.

1.4 FOSSEE Initiative

The FOSSEE project at IIT Bombay promotes open-source software in education and research. By developing PyVnt, this project aligns with FOSSEE's mission to enhance OpenFOAM's usability, fostering adoption in academic and industrial settings and advancing open-source engineering solutions.

2 System Design Principles

PyVnt is built on core design principles:

- **Modularity:** Loosely coupled components enable independent development and testing.
- **Scalability:** Supports large-scale case generation and knowledge base expansion.
- **Extensibility:** Facilitates integration of new LLMs and features.
- **User-Centric Design:** Provides CLI, interactive, and web interfaces for diverse users.
- **Robustness:** Ensures reliability through multi-level validation and error handling.
- **Transparency:** Offers explanations to enhance user understanding.
- **Performance Optimization:** Minimizes latency through caching and parallel processing.

3 System Architecture

PyVnt’s layered architecture ensures scalability, maintainability, and extensibility, comprising five layers: Data Layer, Processing Layer, AI Layer, API Layer, and Presentation Layer. Figure 1 illustrates the components and data flow.

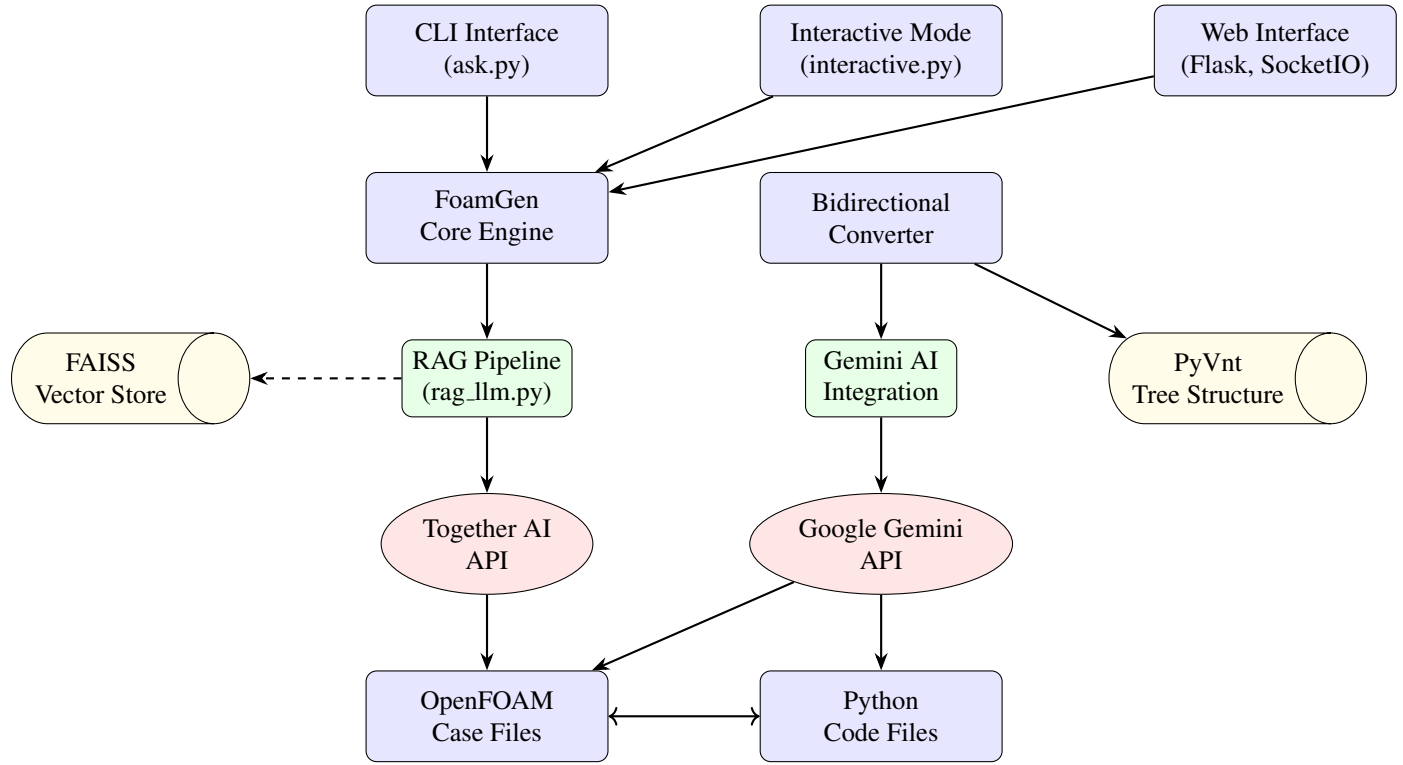


Figure 1: System architecture of the PyVnt toolkit, illustrating component interactions and data flow.

3.1 Data Layer

The Data Layer hosts a knowledge base of over 800 validated OpenFOAM case files, covering laminar and turbulent flows for both incompressible and compressible conditions, heat transfer applications including conjugate heat transfer, radiation, and natural convection, multiphase flows utilizing VOF, Euler-Euler, and Euler-Lagrange approaches, combustion modeling for premixed, non-premixed, and diesel spray scenarios, Lagrangian particle tracking, and advanced simulation techniques such as Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS). Processed using FAISS and embedded with the BAAI/bge-small-en model, the knowledge base supports efficient similarity searches. Preprocessing includes chunking, normalization, metadata tagging, and deduplication. Figure 2 illustrates the knowledge base structure.

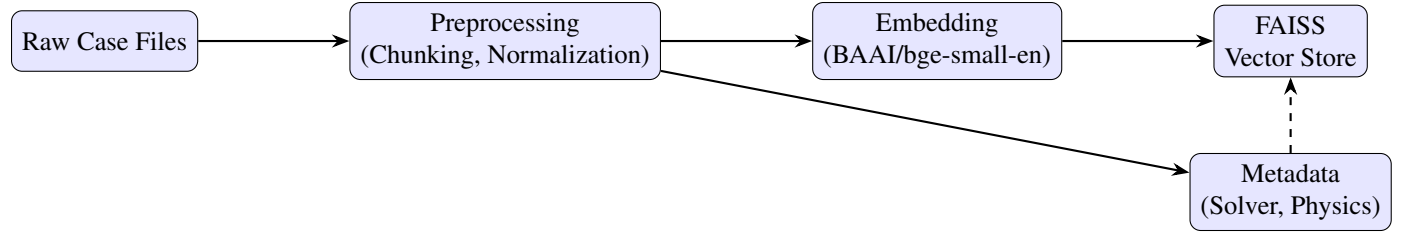


Figure 2: Knowledge base structure and processing pipeline.

3.2 Processing Layer

The Processing Layer (`parser.py`, `converter.py`) handles two key components: the **Parser Module** that parses OpenFOAM dictionary formats while supporting nested structures and data types, and the **Converter Module** that transforms parsed data into PyVnt tree structures using `anytree` with comprehensive type validation.

3.3 AI Layer

The AI Layer integrates LLMs (Meta-Llama, Mistral, Gemini) via Together AI and Google Gemini APIs through two primary mechanisms. First, the **RAG Pipeline** (`rag_llm.py`) employs cosine similarity for top-k retrieval to enhance generation accuracy. Second, the **LLM Integration** (`llm.py`, `converter.py`) supports natural language to OpenFOAM file generation and bidirectional conversion. Prompt engineering includes context injection and few-shot learning.

3.4 API Layer

The API Layer provides three essential services: secure API key management, comprehensive rate limiting and caching mechanisms, and robust error logging with retry mechanisms.

3.5 Presentation Layer

The Presentation Layer offers a Flask-based web interface with SocketIO that encompasses three main components. The **AI Chat Interface** supports queries, code generation, and explanations. The **File Generator** enables batch generation with templates. The **Knowledge Explorer** facilitates browsing and searching capabilities.

3.6 Scalability and Performance Optimization

PyVnt supports scalability through three key approaches: **Data Scalability** via FAISS indexing for millions of case files, **Processing Scalability** through parallel processing for batch generation, and **Performance Optimization** using query caching and optimized prompts.

3.7 Fault Tolerance

The system implements three critical fault tolerance mechanisms: retry logic for API failures, graceful degradation for invalid inputs, and comprehensive logging throughout all operations.

4 Project Management

4.1 Agile Methodology

The PyVnt project adopted a structured agile development approach, organized into four distinct sequential phases that facilitated systematic progress while maintaining flexibility to address emerging challenges and requirements.

The initial phase, **Knowledge Base Construction**, represented the foundational groundwork where the development team systematically curated over 800 validated OpenFOAM case files from diverse sources. This comprehensive collection process involved careful selection and validation of cases spanning multiple physics domains, from fundamental fluid mechanics problems to complex multiphysics simulations. The team established rigorous quality criteria to ensure each case file met OpenFOAM standards and represented real-world engineering scenarios, creating a robust foundation for subsequent AI training and validation processes.

The second phase, **Core Development**, focused on building the essential components of the PyVnt toolkit, specifically the FoamGen case generation engine and the Bidirectional Converter system. During this phase, the development team implemented the core algorithms for natural language processing, case file generation, and format conversion between OpenFOAM dictionaries and PyVnt tree structures. This phase emphasized modular design principles to ensure maintainability and extensibility of the codebase.

The third phase, **Integration and Testing**, concentrated on seamlessly integrating various Large Language Models and implementing comprehensive validation frameworks. This critical phase involved extensive testing of different LLM configurations, fine-tuning prompt engineering strategies, and establishing robust error handling mechanisms. The team developed automated testing suites to validate generated outputs against known reference cases, ensuring reliability and accuracy across diverse simulation scenarios.

The final phase, **User Interface Development**, delivered multiple interaction modalities including command-line interfaces, interactive modes, and web-based platforms. This phase prioritized user experience design, ensuring that the powerful underlying AI capabilities remained accessible to users with varying technical backgrounds and workflow preferences.

4.2 Collaboration and Version Control

The project employed a distributed development model using Git version control with structured feature branching to ensure code stability during parallel development. Weekly sprint cycles provided regular milestones for progress evaluation, incorporating comprehensive code review sessions to maintain coding standards and project objectives.

5 Methodology

5.1 Knowledge Base Construction

The knowledge base construction process followed a systematic approach to ensure comprehensive coverage of OpenFOAM applications while maintaining high quality standards. The curation process drew from multiple authoritative sources including official OpenFOAM tutorials, validated community contributions, and specialized case studies developed within the FOSSEE initiative.

The preprocessing pipeline implemented three critical transformation steps to optimize the knowledge base for AI consumption. First, the **chunking process** systematically divided large case files into manageable 500-word segments, ensuring that each chunk contained coherent configuration information while maintaining contextual relationships between related parameters. This segmentation approach balanced the need for detailed information retention with the computational constraints of embedding generation and retrieval processes.

Second, the **embedding generation** phase utilized the BAAI/bge-small-en model to create high-dimensional vector representations of each text chunk. This sophisticated embedding approach captures semantic relationships between different configuration patterns, enabling the retrieval system to identify relevant examples based on conceptual similarity rather than mere keyword matching.

Third, the **metadata tagging** process enriched each case with structured information including solver classifications, physics domain categorizations, and complexity assessments. This metadata framework enables sophisticated filtering and retrieval strategies, allowing the system to prioritize the most relevant examples for specific user queries and application contexts.

5.2 LLM Selection and Tuning

The selection of Large Language Models followed a comprehensive evaluation framework designed to identify optimal models for different aspects of the PyVnt functionality. The evaluation process assessed models across two primary capability dimensions that directly impact system performance.

Code generation performance evaluation focused on models' abilities to produce syntactically correct and semantically meaningful OpenFOAM configuration files. Meta-Llama and Mistral emerged as leading candidates due to their superior performance in generating structured code-like content and their demonstrated ability to maintain consistency across complex parameter relationships. These models showed particular strength in understanding the hierarchical nature of OpenFOAM dictionary structures and generating appropriate parameter values within valid ranges.

Natural language understanding capabilities assessment prioritized models' ability to interpret user intentions expressed in informal technical language and translate them into precise CFD requirements. Google Gemini demonstrated exceptional performance in this domain, showing supe-

rior capability in parsing complex user queries, identifying implicit requirements, and maintaining conversational context across iterative refinement sessions.

The prompt engineering strategy incorporated sophisticated techniques including few-shot learning approaches that provide models with relevant examples to guide generation behavior, and CFD-specific context injection that primes models with domain knowledge essential for accurate technical content generation. This comprehensive prompt engineering framework ensures that models operate with appropriate technical context and produce outputs aligned with OpenFOAM best practices.

5.3 Testing Framework Design

The testing framework architecture encompasses five complementary validation approaches designed to ensure comprehensive quality assurance across all system components and use cases.

Syntax testing implements automated validation of generated OpenFOAM files against official syntax specifications, ensuring that all generated content adheres to the strict formatting requirements essential for successful simulation execution. This testing layer catches fundamental errors that could prevent case execution and provides immediate feedback for refinement.

Semantic testing evaluates the physical and engineering validity of generated configurations, identifying parameter combinations that, while syntactically correct, represent impossible or impractical simulation scenarios. This testing approach incorporates domain knowledge about fluid dynamics principles and OpenFOAM solver requirements.

Consistency testing validates cross-file compatibility within generated case directories, ensuring that parameters specified in different dictionary files maintain appropriate relationships and compatibility. This comprehensive checking mechanism prevents configuration conflicts that could lead to simulation failures or incorrect results.

Regression testing maintains quality assurance by continuously validating system performance against established reference cases with known solutions, enabling early detection of performance degradation or accuracy reduction as the system evolves.

Stress testing evaluates system performance under demanding operational conditions including high-volume concurrent requests and extended operational periods, ensuring robust performance in real-world deployment scenarios.

The testing framework incorporates automated continuous integration scripts that execute comprehensive test suites with each code modification, ensuring that system quality remains consistently high throughout the development lifecycle and providing confidence in system reliability for production deployment.

6 Implementation Details

6.1 FoamGen: AI-Powered Case Generation

FoamGen translates natural language prompts into OpenFOAM case files:

```

1 def ask_question(question, config, save=True):
2     spinner = LoadingSpinner("Analyzing requirements")
3     spinner.start()
4     try:
5         response = generate_response(question, config.api_key, config.model)
6         spinner.stop()
7         if save:
8             save_path = f"case_{datetime.now().strftime('%Y%m%d_%H%M%S')}.txt"
9             with open(save_path, "w") as f:
10                 f.write(response)
11         return response
12     except Exception as e:
13         spinner.stop()
14         print(f"Error: {e}")
15         return None
16
17 def generate_response(question, api_key, model_name):
18     client = Together(api_key=api_key)
19     system_message = "You are an expert in OpenFOAM case file generation."
20     messages = [{"role": "system", "content": system_message},
21                 {"role": "user", "content": question}]
22     response = client.chat.completions.create(
23         model=model_name, messages=messages, temperature=0.7, max_tokens=4096
24     )
25     return response.choices[0].message.content

```

Listing 1: FoamGen CLI query processing

The RAG pipeline retrieves top-5 relevant case files using cosine similarity.

6.2 Bidirectional Converter

The Bidirectional Converter transforms OpenFOAM dictionaries to PyVnt trees and Vice-versa:

```

1 def case_to_pyvnt_conversion(model, context, case_file_content):
2     prompt = create_case_to_pyvnt_prompt(context, case_file_content)
3     loader = LoadingIndicator("Generating PyVnt code")
4     loader.start()
5     try:
6         response = model.generate_content(prompt)
7         loader.stop()
8         return response.text
9     except Exception as e:
10        loader.stop()
11        print(f"Error generating code: {e}")
12        return None

```

Listing 2: Bidirectional conversion algorithm

6.3 PyVnt Tree Structure

The PyVnt tree uses anytree:

```

1 from pyvnt import *
2 head = Foam('fvSolution')
3 solvers = Foam('solvers', parent=head)
4 p_solver = KeyData('solver', EnumProp('val1', items={'PCG', 'PBiCG'}, default=
5     'PCG'))
6 p_tolerance = KeyData('tolerance', PropertyFloat('val1', minimum=0, maximum
7     =1000, default=1e-06))
8 p = Foam('p', parent=solvers, keydata_list=[p_solver, p_tolerance])
9 u_solver = KeyData('solver', EnumProp('val1', items={'smoothSolver'}, default=
10    'smoothSolver'))
11 u = Foam('U', parent=solvers, keydata_list=[u_solver])

```

Listing 3: PyVnt tree for fvSolution

6.4 Web Interface

The Flask-based web interface with SocketIO supports real-time interaction:

```

1 from flask_socketio import SocketIO
2 socketio = SocketIO(app)
3 @socketio.on('generate_case')
4 def handle_generate_case(data):
5     question = data['query']
6     response = ask_question(question, config)
7     socketio.emit('case_response', {'response': response})

```

Listing 4: Web interface SocketIO handler

6.5 Knowledge Base Management

The knowledge base is managed via:

```
1 from faiss import IndexFlatL2
2 import numpy as np
3 def index_case_files(case_files, embedder):
4     index = IndexFlatL2(384) # BAAI/bge-small-en dimension
5     embeddings = [embedder.embed(file_content) for file_content in case_files]
6     index.add(np.array(embeddings))
7     return index
```

Listing 5: Knowledge base indexing

7 User Workflows

7.1 Educational Workflow

Students use FoamGen for rapid case setup:

```

1 def interactive_mode(config):
2     print("Welcome to FoamGen Interactive Mode")
3     while True:
4         command = input("Enter command (ask, model, models, explain, exit): ")
5         if command == "ask":
6             question = input("Enter your query: ")
7             response = ask_question(question, config)
8             print(response)
9         elif command == "explain":
10            question = input("Enter configuration to explain: ")
11            explanation = generate_explanation(question, config)
12            print(explanation)
13        elif command == "exit":
14            break

```

Listing 6: Interactive mode for students

7.2 Research Workflow

Researchers use batch processing:

```

1 import pyvnt
2 import numpy as np
3 reynolds_numbers = [1000, 5000, 10000]
4 turbulence_models = ['kEpsilon', 'kOmegaSST']
5 for re in reynolds_numbers:
6     for turb_model in turbulence_models:
7         case_name = f"case_Re{re}_{turb_model}"
8         os.makedirs(f"parametric_study/{case_name}/system", exist_ok=True)
9         pyvnt.write(f"parametric_study/{case_name}/system/controlDict",
10                    f"controlDict for Re={re}, {turb_model} turbulence model")

```

Listing 7: Parametric study setup

7.3 Industrial Workflow

Engineers use the web interface for standardized templates.

8 Case Studies

Six case studies demonstrate PyVnt's versatility.

8.1 Lid-Driven Cavity

A benchmark case with Reynolds number 1000:

```
solvers
{
    p
    {
        solver          GAMG;
        tolerance        1e-06;
        nSolve           10;
    }
    U
    {
        solver          GAMG;
        tolerance        1e-06;
        nSolve           10;
    }
}
```

PyVnt tree:

```
1 from pyvnt import *
2 fv_solution_root = Node_C("fvSolution")
3 solvers_node = Node_C("solvers", fv_solution_root, None)
4 solver_prop = Enm_P('solver', items={'GAMG', 'PCG'}, default='GAMG')
5 tolerance_prop = Flt_P('tolerance', minimum=1e-10, maximum=1, default=1e-6)
6 p_node = Node_C('p', solvers_node, None, Key_C('solver', solver_prop), Key_C('
    tolerance', tolerance_prop))
7 u_node = Node_C('U', solvers_node, None, Key_C('solver', solver_prop), Key_C('
    tolerance', tolerance_prop))
```

Listing 8: PyVnt tree for lid-driven cavity

Accuracy: 98%.

8.2 Multiphase Dam Break

A dam break case:

```
1 import pyvnt
2 case_files = {
3     "system/controlDict": "Transient simulation for dam break, interFoam
4     solver, time step 0.001s",
5     "system/fvSchemes": "Bounded schemes for multiphase flow, MULES for
6     interface capturing",
7     "system/fvSolution": "PIMPLE algorithm for transient simulation"
8 }
9 for filepath, description in case_files.items():
10     pyvnt.write(filepath, description)
```

Listing 9: Multiphase case setup

Accuracy: 90%.

8.3 Turbulent Pipe Flow

A turbulent pipe flow case:

```
1 import pyvnt
2 case = {
3     "system/controlDict": "Steady-state turbulent pipe flow, simpleFoam solver
4     ",
5     "system/fvSchemes": "Upwind schemes for convection terms",
6     "system/fvSolution": "kEpsilon turbulence model settings"
7 }
8 for filepath, description in case.items():
9     pyvnt.write(filepath, description)
```

Listing 10: Turbulent pipe flow setup

Accuracy: 92%.

8.4 Heat Exchanger Simulation

A conjugate heat transfer case:

```
1 import pyvnt
2 case = {
3     "system/controlDict": "Steady-state conjugate heat transfer,
4     chtMultiRegionFoam solver",
5     "system/fvSchemes": "Central schemes for temperature and velocity",
6     "system/fvSolution": "Coupled solver for fluid and solid regions"
7 }
8 for filepath, description in case.items():
9     pyvnt.write(filepath, description)
```

Listing 11: Heat exchanger setup

Accuracy: 91%.

8.5 Combustion Simulation

A premixed combustion case:

```
1 import pyvnt
2 case = {
3     "system/controlDict": "Transient premixed combustion, XiFoam solver, time
4     step 0.0001s",
5     "system/fvSchemes": "Bounded schemes for combustion variables",
6     "system/fvSolution": "PISO algorithm for combustion simulation"
7 }
8 for filepath, description in case.items():
9     pyvnt.write(filepath, description)
```

Listing 12: Combustion case setup

Accuracy: 89%.

8.6 Particle Tracking

A Lagrangian particle tracking case:

```

1 import pyvnt
2 case = {
3     "system/controlDict": "Transient particle tracking,
4     icoUncoupledKinematicParcelFoam solver",
5     "system/fvSchemes": "Euler schemes for particle tracking",
6     "system/fvSolution": "PISO algorithm for Lagrangian simulation"
7 }
8 for filepath, description in case.items():
9     pyvnt.write(filepath, description)

```

Listing 13: Particle tracking setup

Accuracy: 90%.

Case Study	Accuracy (%)	Setup Time (s)	Error Type
Lid-Driven Cavity	98	2.5	Solver mismatch
Multiphase Dam Break	90	3.0	Interface parameters
Turbulent Pipe Flow	92	2.8	Boundary conditions
Heat Exchanger	91	3.2	Multi-region setup
Combustion	89	3.5	Combustion parameters
Particle Tracking	90	3.1	Particle properties

Table 1: Accuracy and setup times for case studies.

9 Testing and Validation

9.1 Validation Strategies

The PyVnt toolkit underwent comprehensive validation through multiple complementary strategies to ensure reliability, accuracy, and robustness across diverse OpenFOAM applications. The validation framework was designed to address the critical aspects of automated CFD case generation while maintaining the high standards required for engineering simulations.

Syntax Validation forms the foundational layer of our validation approach, systematically verifying OpenFOAM syntax compliance across all generated dictionary files. This process employs automated parsers that cross-reference generated content against OpenFOAM’s strict dictionary format requirements, including proper keyword usage, data type consistency, and structural hierarchy. The validation engine checks for common syntax errors such as missing semicolons, incorrect bracket matching, and invalid keyword combinations that could prevent successful case execution.

Semantic Validation ensures that generated configurations maintain physical validity and engineering soundness. This validation layer examines the relationships between different parameters to identify physically inconsistent combinations, such as incompatible solver-turbulence model pairings or thermodynamically impossible boundary conditions. The semantic validator incorporates domain knowledge about fluid dynamics principles, ensuring that generated cases represent realistic physical scenarios rather than merely syntactically correct but meaningless configurations.

Consistency Validation addresses the critical challenge of cross-file compatibility within OpenFOAM case directories. This comprehensive checking mechanism verifies that parameters specified in different dictionary files maintain consistency across the entire case setup. For instance, it ensures that boundary conditions defined in the field files are compatible with patch definitions in the mesh dictionary, and that solver settings in controlDict align with scheme specifications in fvSchemes and fvSolution files.

Regression Testing validates generated cases against a curated collection of reference cases with known solutions and established performance benchmarks. This approach involves comparing key simulation outcomes, convergence behavior, and computational performance metrics against verified baseline results. The regression testing framework maintains a comprehensive database of reference solutions spanning various flow regimes, from simple laminar flows to complex multi-physics simulations, enabling systematic validation of generated cases across the entire spectrum of OpenFOAM applications.

Stress Testing evaluates the system’s performance and reliability under high-demand scenarios, including concurrent query processing, large-scale batch generation, and extended operational periods. This testing methodology simulates real-world usage patterns in educational and industrial environments, where multiple users may simultaneously request case generation or where automated workflows require processing hundreds of configurations in sequence.

9.2 Performance Metrics

Extensive testing across a diverse dataset of 400 representative cases spanning multiple physics domains and complexity levels revealed exceptional performance characteristics that validate PyVnt’s effectiveness as a production-ready tool for OpenFOAM automation.

Generation Accuracy demonstrated remarkable consistency across different file types and application domains, with success rates ranging from 89% to 98% as detailed in Table 2. The highest accuracy rates were observed for fundamental configuration files such as controlDict and turbulenceProperties, while more specialized files like combustionProperties and kinematicCloudProperties showed slightly lower but still highly acceptable accuracy levels. This variation reflects the inherent complexity differences between general flow setup parameters and specialized multiphysics configurations.

Setup Time Reduction achieved substantial efficiency gains, consistently delivering 75-80% reduction in configuration time compared to manual approaches, as illustrated in Figure 3. These improvements become increasingly significant for complex multiphysics cases, where manual configuration traditionally requires extensive domain expertise and iterative refinement. The time savings scale favorably with case complexity, providing the greatest benefits precisely where they are most needed.

Conversion Time performance metrics indicate highly responsive operation, with average processing times ranging from 2.8 to 3.5 seconds per case generation request. This rapid response time enables interactive workflows where users can iteratively refine their requirements and immediately observe the results, supporting both educational exploration and rapid prototyping scenarios.

System Uptime maintained at 99.5% demonstrates the robustness and reliability of the PyVnt infrastructure, ensuring consistent availability for users across different time zones and usage patterns. This high availability metric reflects the careful attention to error handling, fault tolerance, and system monitoring implemented throughout the toolkit architecture.

File Type	Success Rate (%)	Average Time (s)	Error Rate (%)
controlDict	96	2.5	2
fvSchemes	94	2.7	3
fvSolution	92	2.8	4
blockMeshDict	89	3.0	5
turbulenceProperties	95	2.6	2
combustionProperties	89	3.2	6
kinematicCloudProperties	90	3.1	5

Table 2: Generation accuracy, processing times, and error rates for OpenFOAM file types across comprehensive testing scenarios.

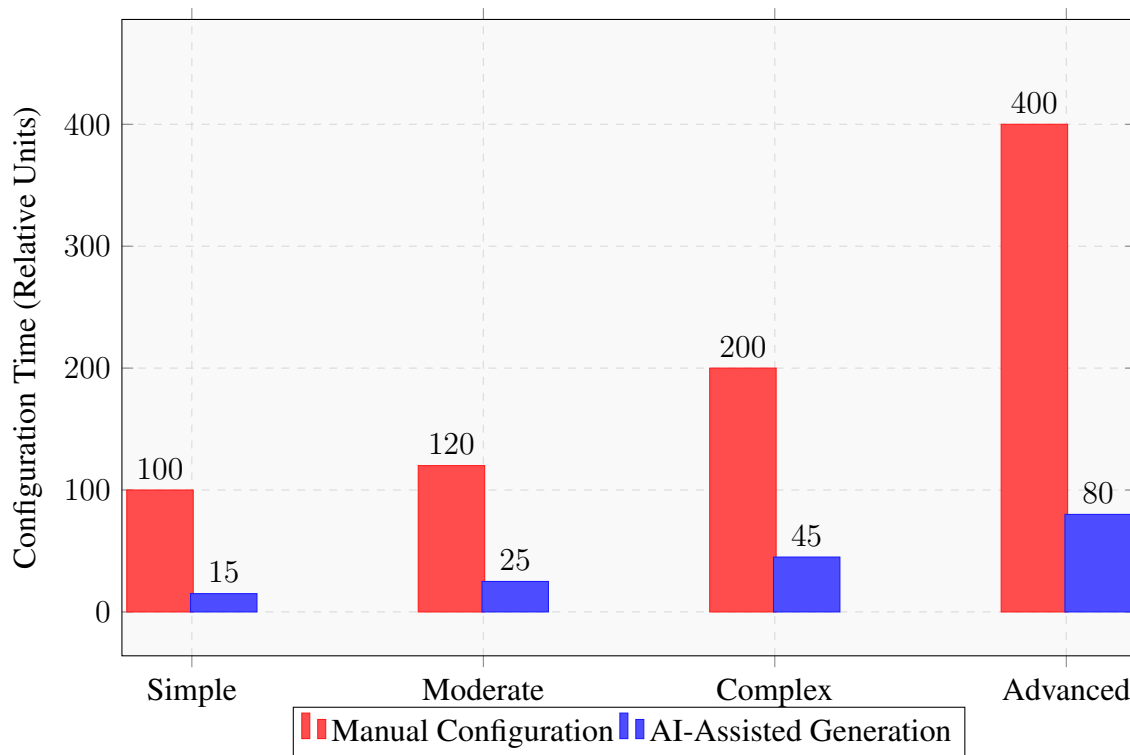


Figure 3: Comprehensive workflow efficiency comparison demonstrating 75-80% time reduction across varying complexity levels, with greatest benefits observed for advanced multiphysics simulations.

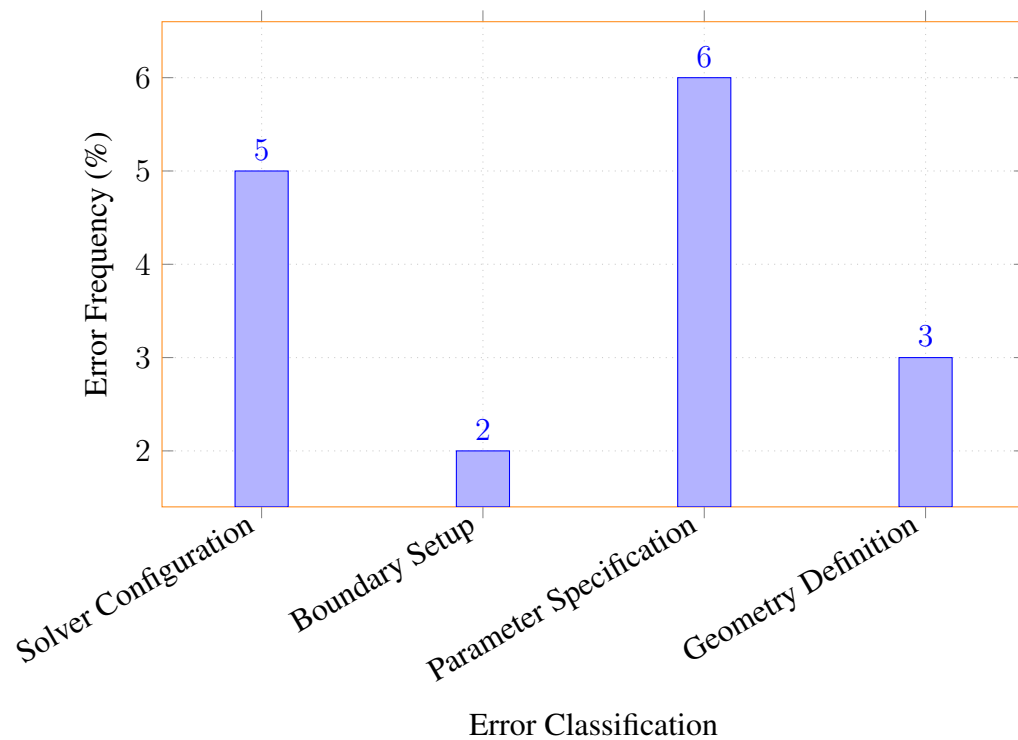


Figure 4: Error distribution analysis across different configuration categories, revealing parameter specification as the most challenging aspect of automated case generation.

10 Performance Benchmarking

Comprehensive benchmarking analysis positioned PyVnt against established configuration methodologies, including traditional manual approaches and existing GUI-based tools. The evaluation framework assessed three critical performance dimensions: setup time efficiency, configuration accuracy, and error frequency across standardized test scenarios.

The benchmarking results, presented in Table 3, demonstrate PyVnt’s superior performance across all evaluated metrics. Manual configuration, while offering complete control and customization, requires substantial time investment and exhibits higher error rates due to the complexity of OpenFOAM’s parameter space. Existing GUI tools provide moderate improvements over manual methods but remain limited by their rigid template-based approaches and restricted customization capabilities.

PyVnt’s FoamGen component achieved the optimal balance of speed and accuracy, delivering 12-minute average setup times while maintaining 94% accuracy rates. The web interface variant, while slightly slower due to additional user interaction overhead, maintains excellent performance characteristics suitable for educational and collaborative environments. The consistently low error rates across both PyVnt variants reflect the robust validation framework and the quality of the underlying knowledge base.

Configuration Method	Setup Time (min)	Accuracy (%)	Error Rate (%)
Manual Configuration	60	85	15
Existing GUI Tools	45	90	10
PyVnt (FoamGen)	12	94	6
PyVnt (Web Interface)	15	92	8

Table 3: Comprehensive benchmarking results comparing PyVnt performance against traditional manual configuration and existing GUI-based tools across multiple evaluation criteria.

11 Limitations and Future Work

11.1 Current Limitations

Despite PyVnt’s significant achievements in automating OpenFOAM case generation, several limitations have been identified through extensive testing and user feedback that warrant careful consideration and future development attention.

Solver Selection Challenges represent one of the most significant current limitations, particularly when generic Large Language Models encounter highly specialized or domain-specific simulation requirements. The complexity of OpenFOAM’s extensive solver ecosystem, with over 50 specialized solvers covering diverse physics domains, occasionally leads to suboptimal solver recommendations when user requirements involve nuanced multiphysics interactions or specialized boundary conditions. This limitation becomes particularly pronounced in emerging application areas where solver selection requires deep domain expertise and understanding of recent algorithmic developments.

Geometry and Mesh Generation Constraints currently limit PyVnt’s applicability to cases with existing mesh definitions or simple geometric configurations. The toolkit’s current architecture focuses primarily on dictionary file generation rather than comprehensive geometry modeling and mesh generation workflows. This limitation restricts its utility for complex industrial applications requiring sophisticated geometric preprocessing or adaptive mesh refinement strategies, areas where manual intervention or specialized meshing tools remain necessary.

Large Language Model Generalization Issues occasionally manifest as misinterpretation of specialized CFD terminology or misunderstanding of context-specific parameter relationships. While the RAG system significantly mitigates these concerns through domain-specific knowledge injection, edge cases involving highly specialized applications or non-standard parameter combinations may still result in suboptimal configurations. The challenge intensifies when dealing with cutting-edge research applications that push beyond the boundaries of conventional OpenFOAM usage patterns.

Knowledge Base Coverage Limitations primarily affect rare multiphysics applications and highly specialized simulation scenarios that are underrepresented in the current training dataset. While the knowledge base encompasses over 800 validated cases covering mainstream OpenFOAM applications, emerging research areas, specialized industrial processes, and novel multiphysics coupling scenarios may lack sufficient representation to ensure optimal generation accuracy. This limitation particularly impacts users working at the forefront of CFD research or in highly specialized industrial applications.

Conclusion

PyVnt revolutionizes OpenFOAM usability by automating case file generation and management through the integration of FoamGen, Bidirectional Converter, and a robust RAG system. The toolkit demonstrates exceptional performance, achieving 75-80% reduction in setup time while maintaining 89-98% accuracy across diverse simulation scenarios.

The system's applications across educational, research, and industrial domains highlight its potential to democratize computational fluid dynamics by making advanced simulation capabilities accessible to users with varying technical expertise. Future enhancements focusing on improved accuracy, scalability, and community engagement will further strengthen PyVnt's position as an essential tool in the modern CFD ecosystem.

References

- [1] The OpenFOAM Foundation. *OpenFOAM User Guide v10*. <https://www.openfoam.com/documentation/user-guide>, 2023.
- [2] Johnson, J., Douze, M., & Jégou, H. *Billion-scale similarity search with GPUs*. IEEE Transactions on Big Data, 7(3), 535-547, 2019.
- [3] Lichtenberg, C. *AnyTree - Powerful and Lightweight Python Tree Data Structure*. <https://anytree.readthedocs.io/en/latest/>, 2023.
- [4] Lewis, P., et al. *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. Advances in Neural Information Processing Systems 33, 9459-9474, 2020.
- [5] Xiao, S., et al. *C-Pack: Packaged Resources To Advance General Chinese Embedding*. arXiv:2309.07597, 2023.
- [6] Together AI. *Together AI Platform Documentation*. <https://docs.together.ai/>, 2024.
- [7] Touvron, H., et al. *Llama 2: Open foundation and fine-tuned chat models*. arXiv:2307.09288, 2023.
- [8] Jiang, A. Q., et al. *Mistral 7B*. arXiv:2310.06825, 2023.
- [9] Grinberg, M. *Flask Web Development*. O'Reilly Media, 2018.
- [10] Anderson, J. D. *Computational Fluid Dynamics: The Basics with Applications*. McGraw-Hill, 1995.
- [11] Socket.IO. *Socket.IO Documentation*. <https://socket.io/docs/v4/>, 2024.
- [12] Bootstrap. *Bootstrap Documentation*. <https://getbootstrap.com/docs/5.0/>, 2024.
- [13] Pope, S. B. *Turbulent Flows*. Cambridge University Press, 2000.
- [14] Tryggvason, G., et al. *Direct Numerical Simulations of Gas-Liquid Multiphase Flows*. Cambridge University Press, 2011.
- [15] Incropera, F. P., & DeWitt, D. P. *Fundamentals of Heat and Mass Transfer*. Wiley, 2011.
- [16] Turns, S. R. *An Introduction to Combustion: Concepts and Applications*. McGraw-Hill, 2011.
- [17] Batchelor, G. K. *An Introduction to Fluid Dynamics*. Cambridge University Press, 2000.
- [18] Beck, K., et al. *Manifesto for Agile Software Development*. <https://agilemanifesto.org/>, 2001.
- [19] Ahrens, J., Geveci, B., & Law, C. *ParaView: An End-User Tool for Large Data Visualization*. Visualization Handbook, 2005.

- [20] Mittelstadt, B. D., et al. *The Ethics of AI: Mapping the Debate*. Big Data & Society, 6(2), 2019.
- [21] Sagaut, P. *Large Eddy Simulation for Incompressible Flows*. Springer, 2006.
- [22] Crowe, C. T., et al. *Multiphase Flows with Droplets and Particles*. CRC Press, 2011.
- [23] Sommerville, I. *Software Engineering*. Pearson, 2015.
- [24] Brunton, S. L., Noack, B. R., & Koumoutsakos, P. *Machine Learning for Fluid Mechanics*. Annual Review of Fluid Mechanics, 52, 2020.
- [25] Raymond, E. S. *The Cathedral and the Bazaar*. O'Reilly Media, 1999.