FINAL REPORT

THE UNIVERSITY OF SHEFFIELD

Department of Chemical & Biological Engineering

Efficient Simulation of Gas-Liquid Multiphase Flow in Rushton Tanks through Integrated VOF and ROM Methods

Supervisor:

Prof. Joan Cordiner

Author: Hengxin Wu

Apr 17, 2026

Signed:

Contents

Abstract

Multiphase flows in stirred-tank reactors are common in industrial processes such as chemical reactions, crystallization, and fermentation. These flows (gas-liquid, liquid-liquid, and solid-liquid) are governed by fluid properties (e.g., viscosity and density), impeller design, and operating conditions. For example, low-viscosity and low-density fluids generate greater turbulence during high-speed mixing. A better understanding of these factors is key to improving mixing efficiency and product quality. Computational fluid dynamics (CFD) simulations are widely used to analyse flow patterns and mixing. However, high computational demands and modelling limitations limit their effectiveness, highlighting the need for more efficient predictive models. To address these needs, this study reviews recent advances in CFD and reduced-order models (ROMs) for stirred-tank multiphase flows, comparing their advantages and limitations. It then presents a high-fidelity ROM developed using ANSYS and validated against full-order CFD simulations. Results show that ROM accuracy depends on the fidelity of the full-order model and the number of snapshots used for ROM training. Improvements to the full-order model were made to generate a richer snapshot dataset, improving ROM performance. Finally, scale-up experiments confirmed that the optimised ROM can accurately and effectively simulate the multiphase flow behaviour in the stirred tank.

Keywords: Multi-phase flow, Stirred tank, CFD simulation, SVD method, Pod method, Volume of Fluid model, Reduced order model

Section 1. Introduction and background

1.1 Contextual Background

Mixed tanks are widely used in various fields, including chemical engineering, biotechnology, and pharmaceuticals. For example, in the chemical industry, mixing tanks are used to mix chemicals such as acids, bases, solvents, and polymers to produce products such as fertilizers, plastics, and adhesives. In the pharmaceutical industry, mixing tanks are used to mix and homogenize active pharmaceutical ingredients (APIs) and other materials used in pharmaceutical production. In the food and beverage industry, mixing tanks are used to produce products such as sauces, condiments, and beverages. Given the wide range of applications, in-depth research on the performance of mixed tanks is crucial to current industry and a key research topic.

The performance of mixed tanks depends on many factors, such as impeller design, fluid properties, operating conditions, and flow regimes. In particular, the dynamics of the multiphase flow within a mixed tank influence mass and heat transfer rates, mixing efficiency, power consumption, and product quality. Therefore, understanding and predicting the flow regime within a mixed tank is crucial for optimising its design, operation, and control.

1.2 Computational Fluid Dynamics in Multi-phase Flow: An Overview

Flows within stirred tanks can be categorised based on the properties of the fluid and its interaction with the impeller. For example, flows can be classified as homogeneous or heterogeneous based on composition or properties (phase distribution, field distribution, etc.); alternatively, flows can be categorised as laminar or turbulent based

on the flow regime (typically based on the Reynolds number). Each flow regime has its own unique characteristics, requiring different modelling and numerical methods for simulation and analysis.

Computational fluid dynamics (CFD) is a key tool for analysing multiphase flows within stirred tanks in many studies because it can efficiently solve the governing fluid dynamics equations (such as the Navier-Stokes equations) to predict flow patterns, mixing behaviour, and other key characteristics within the stirred tank. This predictive capability reduces the economic cost of design exploration while providing valuable and deeper insights into complex fluid dynamics.

The widespread use of commercial and open-source CFD software (such as ANSYS and Open FOAM) has further solidified CFD's position as the most widely used method for simulating and analysing multiphase flows within stirred tanks. However, the high complexity and computational cost of CFD simulations, particularly for turbulent multiphase flows, limit their practicality in large-scale industrial applications. In recent years, reduced-order models (ROMs) have been proposed as a promising approach to overcome these limitations.

This limitation is particularly acute in industrial-scale applications, where experimental data are often scarce due to the prohibitive cost and complexity of large-scale measurements. ROMs offer a pathway to bridge this gap by providing fast, yet reasonably accurate predictions that can support design and scale-up decisions.

1.3 Introduction to Reduced Order Model and VOF Model in Multi-phase Flow

Reduced-order models (ROMs) offer significant advantages over traditional computational fluid dynamics techniques in multiphase flow simulations because they effectively capture the fundamental behaviour of a system while significantly reducing

computational costs. ROMs generated using this technique enable real-time or near-real-time analysis with faster simulation speeds, which is particularly beneficial for process control applications. Traditional CFD remains the preferred method for achieving high-fidelity simulations. For example, Khopkar et al. were among the first to apply modern CFD to a Rushton-turbine vessel with experimental validation, finding good agreement in overall flow structures despite some quantitative discrepancies (1). Their work confirmed that CFD can reproduce key gas-liquid flow features and save the cost of building full-scale prototypes building full-scale prototypes, ROMs offer an excellent balance between accuracy and computational efficiency, making them particularly advantageous for studying multiphase systems under a wide range of scenarios.

Developing ROMs requires snapshot data from high-fidelity, full-order simulations; therefore, a suitable full-scale CFD model must be established to generate the necessary data. This study employed the volume of fluid (VOF) method to construct a full-order model. This method tracks interfaces between fluid phases by assigning volume fractions to each phase (e.g., gas and liquid) within each computational cell, accurately capturing the location and dynamic behaviour of the interfaces.

The VOF method is widely recognised as an efficient and accurate method for simulating gas-liquid multiphase flows and has been validated through numerous validation studies. For example, Mahmud et al. used VOF to simulate the free surface flow caused by magnetic stirring in an unbaffled stirred tank, and the results showed that the predicted flow pattern was highly consistent with the experimental measurements (6). Similarly, Serra et al. showed that the VOF model can accurately reproduce the free liquid surface height fluctuations observed in the experiment when simulating gas-liquid two-phase flow in a stirred tank with a free surface (2). In addition

to using VOF independently, this method can also be combined with other computational techniques to expand its application range. Kang et al. coupled VOF with the discrete element method (DEM) to study the solid-liquid two-phase flow in an unbaffled stirred tank with a free surface. The results showed that the DEM-VOF coupled model can accurately capture the dynamic characteristics of multiphase flow (3). Wu et al. also used a similar VOF model to simulate turbulent gas-liquid flow in a stirred tank with a free surface, further verifying that the model can accurately capture the dynamic characteristics of multiphase flow (4). In addition, Kang et al. used the DEM-VOF model to study the sinking process of floating particles on the free surface of solid-liquid multiphase flow, further confirming that the coupled method has high accuracy and reliability in capturing complex flow behaviour and particle sinking mechanism (5). These studies indicate that the VOF-based full-order model provides a practical framework for simulating gas–liquid multiphase flows in stirred tanks.

The method can be readily coupled with turbulence closures (e.g., RANS models) and particle-tracking approaches (e.g., DEM), which facilitates its application to complex multiphase configurations. However, VOF inherently relies on interface compression schemes and finite grid resolution, which may introduce numerical diffusion or curvature errors in highly turbulent regions.

In the present study, these limitations were mitigated by applying mesh refinement near the free surface and impeller discharge region, together with bounded high-resolution discretization schemes and careful timestep control. Under these conditions, the VOF model was verified against experimental data and demonstrated sufficient accuracy for free-surface prediction within the investigated operating range.

1.4 Objectives and Applications of ROM

The primary goal of reduced-order modelling (ROM) is to mitigate the high computational cost of high-fidelity, full-scale simulations. This approach achieves this by constructing a low-dimensional surrogate model that approximates full-order simulations while preserving the essential flow physics and interface dynamics. This study employed proper orthogonal decomposition (POD) combined with singular value decomposition (SVD) to identify key flow features and remove negligible ones, significantly reducing the system dimensionality and minimising accuracy loss. The resulting ROM significantly reduces simulation time (typically 10 to 100 times faster than a full CFD model) while accurately reproducing detailed flow behaviour. This enables extensive parameter studies and even near-real-time analysis, facilitating design space exploration and rapid design optimisation. Consequently, ROM methods have become a powerful tool in chemical and process engineering, offering a practical balance between accuracy and computational efficiency for high-fidelity multiphase flow simulations.

Reduced-order modelling has recently attracted considerable attention for simulating multiphase flows in stirred tanks (typically using VOF-based CFD) due to its potential to significantly reduce computational cost while maintaining high predictive accuracy. However, developing ROMs that can reliably handle such complex flows remains a major challenge. Ongoing research focuses on improving the accuracy of ROM formulations for strongly nonlinear multiphase flows; effectively integrating ROM frameworks with experimental observations for calibration; and rigorously validating ROM predictions against experimental data to ensure their reliability. Notably, Ahmed et al. provide a comprehensive review of a key aspect of these challenges—evaluating ROM closure strategies by comparison with the neglected modes—covering

approaches from first-principles methods to machine learning techniques, and highlighting the advantages, limitations, and representative applications of each (34). This study provides valuable insights into developing more robust and efficient ROMs for complex multiphase systems. The practical benefits of ROM in industrial contexts can be further clarified. The ROM developed in this study reduces simulation time from days to minutes once trained, enabling rapid exploration of multiple operating conditions. This efficiency is particularly valuable for industrial applications, where experimental data for large-scale tanks are often scarce. By providing fast and reasonably accurate predictions, ROMs support design exploration, scale-up decisions, and integration into digital twin frameworks. For example, a single high-fidelity CFD simulation of the stirred tank at a given operating point typically requires 12–24 hours of computation on a high-performance workstation. In contrast, once trained, the ROM produces comparable results in less than one minute. This speedup allows engineers to evaluate dozens of design variants (e.g., different impeller speeds, eccentricities, or tank scales) within a single day, which would be impractical with full CFD alone. Reduced-order models can significantly reduce the computational burden while preserving the key dynamic characteristics of multiphase flow systems. For example, Gazzani et al. developed a ROM for a Shell-Prenflo entrained flow gasifier by constructing a compartment model, which represents the complex fluid dynamics and chemical reactions using a network of interconnected idealized reactors, such as continuous stirred-tank reactors (CSTRs) and plug flow reactors (PFRs) (48). This physics-based approach successfully captured the main dynamic behaviours of the gasifier and significantly reduced the simulation time. This successful demonstration highlights the great potential of ROM as an efficient modelling strategy for complex multiphase processes. Similarly, ROM technology has been successfully applied to the simulation

of multiphase flows in stirred tank reactors, enabling simulation-based reactor design and operation optimisation. In contrast, conducting similar comprehensive studies using full-scale CFD simulations would be time-consuming and uneconomical. Furthermore, integrating advanced data-driven methods, such as deep learning, into the ROM framework promises to further improve the model's predictive accuracy without sacrificing computational efficiency. Therefore, continued improvement of ROM techniques and extension of these models to new multiphase flow domains remain active areas of research.

1.5 Research Aims and Objectives

The primary goal of this study is to construct a highly accurate reduced-order model (ROM) to develop a new approach to understanding multiphase systems, with a particular focus on its application in industry. To this end, it will optimise key elements of ROM construction, such as the selection of parameter design points (i.e., the specific combinations of input parameters, such as impeller speed and fluid properties, at which high-fidelity simulations are performed to generate training data) and the number of model modes (i.e., the number of retained basis functions in a POD-based ROM), to achieve a balance between computational efficiency and predictive accuracy in complex multiphase flow simulations.

To achieve this goal, this study will systematically develop and optimise the ROM model and apply it to real-world industrial scenarios. Specifically, select appropriate parameter design points, determine the optimal number of reduced-order modes, and thoroughly evaluate the impact of the fidelity of the full-order CFD model on the reliability of the ROM under various operating conditions. Subsequently, use the constructed ROM to predict key output parameters (such as liquid level, static pressure,

and other flow characteristics) in a stirred tank process to validate the effectiveness of this approach in real-world applications.

While the ROM has previously been initially validated in laboratory-scale stirred tanks, this study will extend this evaluation to larger-scale stirred tank systems. The applicability of the ROM under these large-scale conditions will be evaluated through numerical simulations and comparisons with experimental results. This expansion represents a continuation of the current work, seamlessly integrating with existing goals of industrial relevance and model validation, rather than a standalone research objective. Its focus is on validating the robustness of the model under industrial scale-up conditions.

Furthermore, rigorous validation and sensitivity analysis will be performed on both simulation and experimental data to ensure the robustness of the proposed approach. All research results and methods will be fully documented and publicly available, providing valuable insights to stakeholders in academia and industry.

1.6 Structure of the Report

This report is divided into several sections. Section 2 provides a comprehensive literature review, covering traditional CFD methods and the rise of reduced-order modelling (ROM) techniques for multiphase flows in stirred tanks. It discusses the challenges of high-fidelity multiphase CFD, recent advances (e.g., VOF, DEM coupling), and key insights from the literature on reduced-order modelling (including its advantages, applications, and recent developments). Section 3 details the methodology used in this study, including the steps for constructing a ROM. It describes the generation of snapshots from full-order simulations, the application of SVD/POD for pattern extraction, and the ROM construction and validation procedures. Section 4

presents the data sources and experimental case details used in this study. This includes the Rushton tank geometry, material properties, operating conditions, and the CFD model established using the multiple reference frame (MRF) method. Section 5 presents simulation results and analysis. This study examines the behaviour of key output parameters (e.g., level distribution and pressure field) and compares the performance of ROM predictions with full CFD results at various design points. Section 6 explores strategies for improving ROM accuracy. This section analyses the sources of error in the ROM and proposes specific improvement measures—for example, adjusting turbulence model considerations, optimising viscosity model parameters, and examining the impact of key model coefficients (Cmu, C1, C2, etc.) on the ROM results. Section 7 evaluates the robustness of the developed ROM and full-order model through additional analyses. This includes validation measures such as grid independence tests and convergence checks, as well as a discussion of the applicability of the ROM beyond the training scenario. Finally, Section 8 summarises the main findings and contributions of the report. It also outlines suggestions for future work, such as extending the ROM method to different tank configurations or incorporating advanced data-driven techniques to further enhance predictive capabilities.

1.7 Research Problem and Scope

This thesis addresses the need for fast yet reliable prediction of gas–liquid free-surface behaviour and bulk circulation in unbaffled Rushton-stirred vessels. In practice, engineers must evaluate multiple operating points (impeller speed, fluid properties) and geometric variants (clearance, impeller diameter, eccentricity) under time and budget constraints. High-fidelity CFD—especially interface-capturing simulations that resolve the free surface—can reproduce the main features of these flows but is computationally

expensive and sensitive to modelling choices; as a result, full parametric studies and uncertainty quantification become impractical. Reduced-order models (ROMs) promise orders-of-magnitude speedups while retaining the dominant flow physics needed for design and control (3).

The central problem is to construct and validate a ROM that accurately predicts: (i) free-surface depression (vortex shape and depth), (ii) circulation topology (primary/secondary cells, recirculation strength), and (iii) key field variables (velocity and pressure distributions) in an unbaffled, gas–liquid Rushton tank over a specified operating envelope. The ROM is trained on snapshots (i.e., high-fidelity CFD solution fields, such as velocity, pressure, and phase fraction distributions, saved at specific operating conditions or time instances) generated by a verified full-order CFD model based on the volume-of-fluid (VOF) method and assessed on held-out conditions and a scale-up case. The goal is to achieve substantial computational savings relative to full CFD while preserving predictive accuracy for vortex-related metrics that are crucial to design and scale-up (4).

System definition and operating envelope: The reference system is a cylindrical, flat-bottomed vessel with a six-bladed Rushton turbine. The study considers Newtonian liquids (e.g., water or DMSO) and an overlying gas phase (air) with a free surface. The operating envelope spans the high-Reynolds-number regime typical of laboratory and pilot scales, where inertial effects dominate and free-surface vortices form. Reported speeds include representative low, medium and high rpm to expose changes in vortex depth and circulation strength with Fr and We. Within this envelope, the ROM is expected to interpolate across speeds and modest geometry variations represented in the training snapshots (12).

Scope and limitations: The CFD training model treats the phases as incompressible and isothermal, with no mass transfer across the free surface and no chemical reaction; phase change, surfactants and bubble breakup/coalescence are out of scope. A single shared-velocity, shared-pressure VOF formulation is used to capture the interface; dispersed bubbly regimes with very high gas volume fraction are not targeted. Turbulence is modelled with RANS closures (standard and improved k–e; k–ω SST; RSM where noted) due to their efficiency; LES is discussed but used sparingly because of cost. The impeller is represented with MRF for mean fields, and sliding-mesh/URANS is used for selected unsteady validation. The numerical schemes use bounded high-resolution schemes; Courant and capillary timestep criteria are enforced; mesh refinement is concentrated near the free surface and impeller discharge. These choices reflect state-of-practice for free-surface tanks and are documented to facilitate replication.

Research questions and hypotheses:

RQ1: How many modes (r) and which snapshot design are required for the ROM to recover free-surface shape and circulation topology with acceptable error across the operating envelope? H1: An r in the low-to-mid teens (e.g., 10–20 modes), selected by combined energy capture and held-out reconstruction error, suffices for accurate vortex-depth and free-surface profile prediction (see, e.g., (2,3) for discussions on mode selection criteria).

RQ2: How sensitive is ROM accuracy to full-order model (FOM) numerical controls (mesh density near the interface, timestep, turbulence constants) and to turbulence model choice?

H2: ROM error mirrors FOM bias; improving FOM verification and calibration reduces ROM error more effectively than increasing r. This is supported by studies that have

demonstrated the sensitivity of ROM accuracy to the fidelity of the underlying full-order simulations (5,10). For example, Khopkar et al. (1) showed that while CFD can capture key gas–liquid flow features, quantitative accuracy depends on model calibration.

RQ3: Does the ROM generalise to modest scale-up and geometric variation (e.g., impeller clearance changes or eccentricity) represented by out-of-sample snapshots?

H3: With targeted snapshots spanning governing-parameter space (Re/Fr/We), the ROM interpolates reliably and preserves qualitative flow topology; extrapolation beyond the training manifold degrades gracefully (4)(11).

Contributions:

- A VOF-based, verified CFD dataset for an unbaffled Rushton tank across multiple speeds and materials, including mesh/timestep sensitivity and calibrated turbulence-model settings (5)(10).
- A POD/SVD ROM that reconstructs free-surface profiles and bulk circulation at a fraction of CFD cost, with explicit criteria for selecting r and for snapshot design across Re/Fr/We (2)(3).
- A validation protocol that compares ROM predictions to both full-order simulations and literature data at an engineering-scale case, emphasizing free-surface metrics relevant to entrainment and scale-up (4)(6)(11).
- A practical workflow for industrial adoption: how to generate training snapshots, verify/validate the FOM, select ROM hyperparameters, and report uncertainty, all within a reproducible pipeline (12)(13).

Section 2. Literature review

2.1 CFD Simulation of Gas–Liquid Stirred Tanks: Progress and Challenges

Multiphase computational fluid dynamics (CFD) has become indispensable for analysing gas–liquid stirred tanks, offering insight into flow patterns, mixing efficiency, and free-surface dynamics without the cost of extensive physical trials. Early CFD studies demonstrated the potential of two-phase models to capture basic flow features, but they faced significant limitations. For example, Serra et al. developed a time-dependent two-phase CFD model for a baffled stirred tank with a free surface, but it was restricted to incompressible liquid–gas flows and struggled to accurately predict turbulent quantities (3). Over the past two decades, rapid advances in CFD algorithms and computing power have improved realism. Khopkar et al. were among the first to apply modern CFD to a Rushton-turbine vessel with experimental validation, finding good agreement in overall flow structures despite some quantitative discrepancies. Their work confirmed that CFD can reproduce key gas–liquid flow features and save the cost of building full-scale prototypes (1). Subsequent researchers addressed CFD's early shortcomings by incorporating more sophisticated physics and numerical techniques. For instance, coupling CFD with particle-tracking has enabled simulations of solid–liquid–gas systems; Kang showed that a coupled discrete element method–Volume of Fluid (DEM–VOF) approach reliably captured liquid–solid hydrodynamics and free-surface behaviour in an unbaffled tank (4). Similarly, Wu introduced a VOF–DEM model for turbulent free-surface flows with suspended particles, successfully predicting complex phenomena like particle dispersion and surface deformation (5). These studies underscore those continual improvements – such as multiphase coupling

and advanced interface tracking – have extended CFD's applicability and accuracy for stirred-tank problems. Despite this progress, simulating gas–liquid stirred tanks remain challenging due to the inherently complex physics. Gas entrainment, bubble dynamics, and impeller-induced turbulence introduce strongly nonlinear behaviours that demand high fidelity. Researchers like Deglon and Meyer emphasise that achieving convergence and accuracy in stirred-tank CFD often requires years of refinement of models and numerical schemes (2)(7). A key area of improvement lies in the modelling of interphase forces. For example, Lane et al. (8) developed a new correlation for the bubble drag coefficient that accounts for the interaction between bubbles and turbulent eddies, demonstrating the importance of incorporating such effects in gas–liquid stirred tank simulations. Likewise, Li et al. (9) systematically evaluated five turbulence models for a magnetic-rod-driven unbaffled mixer and found that the Reynolds Stress Model (RSM) and the SST k–ω model produced significantly better agreement with experimental velocity fields than the standard k–ε model, highlighting the importance of selecting turbulence closures that can capture anisotropy in strongly swirling flows. These findings highlight that while CFD can capture the qualitative flow patterns, precise quantitative predictions require careful selection of models (e.g. for turbulence and multiphase forces) and often case-specific calibration.

In summary, the evolution of CFD for stirred tanks has greatly improved its reliability – from basic two-phase approximations to comprehensive multiphase formulations – yet key challenges persist in capturing turbulence, phase interactions, and scale effects. The following sections discuss in detail the main technical considerations in modelling gas–liquid stirred tanks, as well as emerging methods to alleviate the high computational cost.

2.2 Interface-Capturing Models for Multiphase Flows (VOF, LS, EE)

An accurate representation of the gas–liquid interface is crucial for predicting free-surface vortices and gas entrainment in stirred tanks. Three major approaches are prevalent for interface capturing in reactor-scale simulations: Volume of Fluid (VOF), Level-Set (LS), and Euler–Euler (two-fluid) formulations (14). VOF methods represent the interface by a colour function (volume fraction) that is advected with the flow (16). This approach is inherently mass-conservative and has been widely adopted for gas–liquid free surfaces in mixing vessels (2). The classical VOF technique introduced by Hirt and Nichols (1981) ensures that each computational cell contains a well-defined fraction of gas or liquid, allowing the interface shape to emerge from the contour of this fraction (16). Modern VOF implementations include high-resolution compressive differencing schemes to minimise numerical smearing of the interface. VOF has proven effective when the gas occupies a relatively small portion of the volume and a coherent free surface is present (10). Under such conditions, VOF can sharply capture the deformation of the liquid surface – such as the dip (vortex) that forms under an impeller – as well as waves and sloshing dynamics. Numerous validation studies confirm VOF's accuracy in stirred tanks. For example, Mahmud et al. showed that a VOF simulation of a magnetically driven unbaffled tank reproduced the measured free-surface shape and flow patterns with high fidelity (6, 7). Similarly, more recent VOF-based simulations (often coupled with DEM for solid particles) have been used to optimise tank designs; Kang demonstrated good agreement between DEM–VOF predictions and experimental trends for solids drawdown and surface behaviour (4). However, VOF methods become less effective when the gas phase volume is large or the interface breaks into dispersions. In highly aerated, churn-turbulent regimes with many bubbles,

a single sharp interface no longer exists to be captured by VOF. Under those conditions, VOF without additional sub-grid models cannot resolve the myriad small bubbles and fragmented interfaces, leading to under-prediction of gas hold-up and improper phase topology (10). Even for moderate two-phase flows, a straightforward VOF simulation can fail to capture the time evolution of complex interface shapes, partly due to issues like contact-line pinning and unmodelled 3D instabilities (21). Such limitations must be recognised when applying VOF: interface sharpening techniques and breakup/coalescence models are often needed for highly turbulent gas–liquid flows.

Level-set methods (LS) provide an alternative by advecting a smooth signed-distance function whose zero level corresponds to the interface, inherently yielding a well-defined interface with continuous curvature and normal vectors. These characteristics are advantageous for computing surface tension and handling topology changes such as interface merging or pinching. Osher and Sethian (1988) originally introduced the LS framework for capturing moving interfaces (18). The primary drawback is that LS is not inherently mass-conserving—over time, numerical errors can create or lose fluid volume if no corrective measures are taken. To address this, hybrid VOF–LS schemes have been developed; for example, Sussman and Puckett (2000) combined VOF's conservation with LS's geometrical accuracy, achieving both volume preservation and a smooth interface representation (19). In practice, pure LS is less common in engineering studies of stirred tanks because strict mass conservation is critical for long transients. VOF remains the more popular choice in industrial CFD codes due to its robustness and the simpler implementation of surface tension via the Continuum Surface Force model (Brackbill et al., 1992) (17). Still, LS-based research provides valuable insights and can be beneficial for specific scenarios – for example, predicting

detailed curvature-driven phenomena (e.g. small capillary waves) where VOF's interface may be too diffuse.

Euler–Euler models take a fundamentally different approach by treating the gas and liquid as interpenetrating continuous phases rather than tracking a sharp interface. Each phase has its own set of conservation equations coupled through interphase exchange terms for momentum, mass, and energy (23). Classical two-fluid model formulations by Ishii & Hibiki (2010) (23) and Drew & Passman (1999) (24) introduce drag, lift, virtual mass, and turbulence dispersion forces to represent momentum exchange between bubbles and the liquid phase. EE models do not explicitly resolve a sharp interface; instead, they capture the averaged distribution of phases (e.g., local gas volume fraction). This approach excels in flows with high gas fractions and dispersed bubble fields – conditions typical of vigorous aeration where individual gas pockets are small and numerous. In such churn-turbulent regimes, EE models can reproduce large-scale behaviours like overall gas holdup, circulation patterns, and mixing time trends at a much lower computational cost than interface-resolving methods (22). For instance, van't Riet & Smith's classic experiments (12) and Nienow's correlations (13) established how gas throughput and power input affect aerated mixing; Euler–Euler simulations, when calibrated to those data, have predicted gas holdup and mass transfer rates reasonably well in design studies. Some hybrid approaches address this limitation by, for example, adding a third pseudo-phase to represent a large gas cavity at the free surface while using EE for the dispersed bubbles beneath. Overall, VOF and LS methods are preferred for coherent interface problems (such as predicting the vortex shape in an unbaffled tank at moderate aeration), whereas Euler–Euler is more suitable for dispersed multiphase problems (such as heavily aerated, bubbly flows). In partially aerated industrial vessels, a combination of approaches may even be applied – e.g.,

VOF to capture the main free-surface vortex and EE to model entrained bubbles below. Table 1 provides a brief comparison of these interface modelling strategies adapted from literature (30) and (31), highlighting their strengths, limitations, and typical usage regimes in mixing applications.

In summary, accurate interface modelling in CFD requires choosing a method consistent with the flow regime. Volume of Fluid remains the workhorse for most gas–liquid stirred tank simulations due to its conservation properties and availability in commercial codes. Level-set methods contribute additional mathematical elegance for interface curvature calculations but are mainly used in research. Euler–Euler models offer a pragmatic solution for very high gas fractions where resolving each interface is impractical (29). Contemporary research continues to refine these methods – for example, developing more advanced interface compression schemes for VOF and better closure correlations for two-fluid models – to improve predictive accuracy for complex multiphase reactors (32). The next section discusses another critical aspect of CFD modelling in stirred tanks: turbulence representation and its interplay with multiphase phenomena.

2.3 Turbulence Modelling and Multiphase Interaction in Stirred Tanks

Turbulence is a dominant feature of industrial stirred tanks, especially at the high Reynolds numbers typical of production‑scale equipment. Accurately capturing turbulent mixing and the resulting vortex formation is essential for reliable CFD predictions. The Navier–Stokes equations fully describe turbulence in principle, but solving them directly (DNS) in a large stirred vessel is computationally infeasible. Instead, turbulence models are employed to approximate the effects of eddies on the mean flow. The standard two-equation models, such as the k–epsilon model and its

variants, remain widely used as baselines in stirred tank simulations (27). These models introduce additional transport equations for turbulence kinetic energy (k) and its dissipation rate (epsilon), allowing calculation of an eddy viscosity that closes the mean momentum equations. However, applying such models to swirling, multiphase flows require care. In baffled tanks or those with strong circulation, flow conditions violate some of the k–epsilon model's core assumptions (like isotropy of turbulence). Researchers have therefore tested refined models in this context. For example, Shih et al. developed a realizable k–ε model that adjusts the turbulence scales to remain physically plausible in strained flows (28). Another widely used variant is the RNG k–ε model, originally developed by Yakhot and Orszag (1986) for single-phase flows, which has since been successfully extended to multiphase stirred tank simulations (29, 34). These variants often improve predictions of swirl intensity and recirculation in stirred tanks compared to the standard formulation (34). Another popular choice is the k–omega SST model (Menter, 1994), which blends the strengths of k–omega in the near-wall region with k–epsilon in the far field. The SST model has been recommended for flows with adverse pressure gradients and separation, conditions encountered around impeller blades and baffles (34). In simulations of baffled tanks, k–omega SST often yields better free-surface vortex depth predictions than k–epsilon because of its enhanced sensitivity to vorticity transport in the core flow (34). For cases with pronounced anisotropy – such as strong swirling and the impeller tip vortices – Reynolds Stress Models (RSM) can further enhance fidelity by solving separate transport equations for each stress component, thereby capturing anisotropic turbulence effects (e.g. normal-stress differences) that influence the shape of the air–water interface. For instance, RSM has been used to predict the surface deformation in

unbaffled spinning tanks, achieving closer agreement with experiments than two-equation models in those highly anisotropic flows (29).

In addition to steady RANS approaches, transient turbulence modelling can be important when time-dependent phenomena are of interest. In stirred tanks with periodic blade passage or unstable flow patterns, Unsteady RANS (URANS) and hybrid approaches like Detached Eddy Simulation (DES) or Large Eddy Simulation (LES) have been applied. URANS (with a sliding mesh technique) can capture periodic fluctuations such as impeller blade passage frequencies that a Multiple Reference Frame steady approach would miss. For example, using a sliding mesh URANS approach can reveal transient vortex formation and shedding behind baffles that a frozen-rotor model cannot capture. Fully resolving the turbulence spectrum with LES yields even more detail – such as the chaotic meandering of the vortex and fine-scale eddies – but at substantially higher computational cost. In free-surface flows, LES requires special treatment to avoid excessive interface instability and to model the damping of turbulence at the phase boundary. Rodi noted that standard two-equation models tend to overpredict turbulence in the vicinity of a free surface unless a damping term is included (38). As a result, most RANS-based free-surface simulations include such turbulence damping in cells at the interface to mimic the stabilizing effect of surface tension and stratification. This is another example of model tuning needed specifically for multiphase flows.

Surface tension and contact line effects can be significant in certain regimes, particularly for smaller-scale systems or those with additives that alter interfacial tension. The Continuum Surface Force (CSF) method (Brackbill et al., 1992) is commonly used to model surface tension as a volumetric force acting on the interface in VOF simulations (17). While generally effective, CSF can introduce spurious

“parasite” currents if the computed interface curvature is noisy. Fine meshes and curvature smoothing practices are recommended to mitigate this. When the free surface intersects solid walls (as in partially filled vessels or around shaft or baffle penetrations), proper wall-wetting (contact angle) models are required. Afkhami and Bussmann (2009) implemented static and dynamic contact angle boundary conditions in VOF simulations, showing that these are essential to predict how the liquid clings to or detaches from vessel walls (21). Neglecting contact angle physics can lead to unrealistically pinned or overly mobile contact lines, distorting the simulated free-surface shape. In large industrial tanks, surface tension effects are often weak (high Weber number flows), but in lab-scale or precision applications they can affect bubble entrainment and vortex stability.

Another consideration is the modelling of the impeller’s motion. Two principal approaches exist: the Multiple Reference Frame (MRF) technique and the transient sliding mesh (SM) approach. The MRF method treats the impeller region in a rotating reference frame and the rest of the tank in a stationary frame, coupling them via a momentum source at the interface between the two zones. This yields a steady-state approximation of the mean flow and is computationally economical. MRF has been widely used to predict time-averaged quantities like power draw and circulation patterns in design studies. However, it cannot capture unsteady blade-fluid interactions or any periodic fluctuations. The sliding mesh approach, by contrast, explicitly rotates the impeller geometry through the fluid domain over time, thus resolving the periodic, unsteady wake shedding and turbulence modulation caused by the impeller’s passage. Studies have shown that SM simulations can predict phenomena like cyclic variations in torque, transient gas entrainment by blade tip vortices, and the precession of the vortex core — features not captured by steady-state methods (22,41). For example, in

an aerated tank, a sliding mesh URANS model can capture the intermittent drawdown of gas by the impeller and its effect on the free surface, whereas an MRF model would only give an averaged gas cavity shape. The downside is the much higher computational cost for SM, roughly proportional to the number of time steps required for adequate temporal resolution. For many engineering purposes, MRF provides a reasonable compromise: it predicts the mean vortex depth and power number adequately for scale-up correlations. However, when detailed unsteady flow features or scale-dependent fluctuations are of interest, a sliding mesh or even a hybrid LES approach is warranted to ensure those dynamics are not lost.

Finally, numerical best practices such as grid refinement and time-step control are vital for trustworthy CFD results. Adequate mesh resolution near the free surface and in the impeller region must be ensured to capture steep gradients. High-resolution (bounded) differencing schemes for convection are recommended to prevent unphysical interface diffusion and to maintain stable computations of momentum and volume fraction. Commonly, the SIMPLE or PISO algorithms are used for pressure–velocity coupling in incompressible flow, with convergence acceleration via multi-grid solvers. The time-step in transient simulations is constrained by a Courant number criterion (for advection stability) and by a capillary time scale criterion (to resolve surface oscillations without numerical instabilities) (47). Meeting these criteria often means a very small time-step when fine meshes and high surface tension are involved, which again highlights the computational intensity of high-fidelity CFD. Verification and validation procedures are therefore indispensable to quantify the uncertainty in CFD predictions. Grid independence tests (e.g. using the Grid Convergence Index of Roache, 1998) should be performed, and simulation results should be benchmarked against experimental data such as velocity profiles (e.g. from PIV or LDA) or measured vortex depths (48). In the

context of stirred tanks, validation might include comparing predicted flow numbers, gas hold-up, or mixing time against established empirical correlations and experimental data, such as those reported by Nagata (33) for general mixing parameters and by Ciofalo et al. (24) for unbaffled tank turbulence. By adhering to such best practices, researchers can ensure that their CFD models are both quantitatively reliable and practically useful for design decisions.

In summary, turbulence modelling in stirred-tank CFD requires a judicious choice of closure model and possibly model tuning for multiphase effects. No single turbulence model is universally superior – k–e remains a reasonable starting point, but SST or RSM models provide improvements for swirling flows, and LES/URANS are valuable for unsteady phenomena. The interplay between turbulence and the gas–liquid interface adds further complexity, necessitating special treatments like interface turbulence damping (38). The fidelity of CFD predictions in stirred tanks thus hinges on both physical modelling (turbulence, interphase forces, etc.) and numerical resolution. Even with careful modelling, however, high-fidelity CFD of large tanks is computationally intensive. This motivates the development of reduced-order modelling approaches, as discussed next, to expedite simulations while retaining essential physics.

2.4 Reduced-Order Modelling (ROM) Techniques for Multiphase Flows

Given the expense and complexity of full CFD, reduced-order models (ROMs) have emerged as a promising strategy to simulate stirred-tank flows much more efficiently (51, 52). A ROM is essentially a low-dimensional surrogate of the full hydrodynamic system that preserves the dominant flow characteristics. The central idea is to project the high-resolution CFD data onto a smaller set of basic functions or modes, thereby capturing the key dynamics with far fewer degrees of freedom (35, 59). In multiphase

flows, developing a reliable ROM is challenging due to the strongly nonlinear and transient nature of phenomena like gas entrainment and vortex formation. Nonetheless, recent studies have shown that ROMs can indeed reproduce the main features of gas–liquid stirred tanks at a fraction of the computational cost. The primary benefit of ROMs is computational speed: by simplifying the model equations, ROMs can run orders of magnitude faster than full 3D CFD (53, 55). This enables real-time simulation and extensive parametric studies that would be impractical with CFD alone. For instance, a ROM can be executed quickly to predict how a change in impeller speed or gas flow rate will affect liquid height and circulation, allowing engineers to explore numerous scenarios without rerunning expensive CFD each time (56). ROMs are therefore extremely useful for applications such as process optimisation, control, and what-if analyses in digital twins.

There are several techniques to construct ROMs, but one of the most widely used in fluid dynamics is Proper Orthogonal Decomposition (POD) combined with Galerkin projection (35, 59). In POD, a set of snapshot data from a full CFD simulation (e.g. instantaneous velocity and pressure fields at various times or operating conditions) is decomposed into an orthogonal basis of modes ranked by their energy content. The leading few modes capture the most energetic and coherent structures in the flow (for example, the primary circulation loops, large vortices, etc.) (59). By truncating to these top modes, one can represent the flow state with far fewer variables. The system dynamics can then be projected onto this reduced subspace – often via Galerkin projection, which yields a system of ordinary differential equations governing the time evolution of the modal coefficients (35). This methodology has been successfully applied to stirred tanks. For example, Mikhaylov et al. (2021) used POD to reconstruct large-scale flow structures in a stirred tank from limited sensor data (37). Their POD-

based ROM successfully reconstructed the 3D flow field and achieved a much faster computation than full CFD, demonstrating the potential for real-time state estimation in mixing processes (37). Similarly, Espinoza (38) applied POD and Dynamic Mode Decomposition (DMD) to flow data from stirred tanks and bubble plumes, enabling the identification of dominant flow patterns and reconstruction of velocity fields with a low-dimensional model (65). Other studies have further explored the application of ROM techniques in stirred tanks (63, 64). These examples underline that ROM techniques can maintain high predictive accuracy for key flow features – such as circulation strength, vortex shape, and oscillation frequencies – while reducing simulation times from days to hours or less.

In addition to classical POD-Galerkin ROMs (often called intrusive ROMs because they derive reduced equations directly from the full governing equations), there are also non-intrusive ROMs based on data-driven approaches (36,39). These use machine learning or regression techniques to learn the mapping from inputs to outputs (or the evolution of modal coordinates) without explicitly projecting the Navier–Stokes equations (34). By training on high-fidelity simulation data, an AI-based ROM can potentially infer flow behaviour in real time without solving the full CFD equations, demonstrating the power of combining reduced-order modelling with modern machine learning (36, 40).

In the context of stirred tanks, researchers have begun integrating ROMs with advanced surrogate modelling methods, including deep learning frameworks (39, 41, 44). Such methods – for example, using neural networks to capture the parametric dependence of POD mode coefficients – can automatically handle nonlinearity and complex parameter interactions better than traditional ROM approaches. For instance, Pant et al. (45) developed a DL-ROM (Deep Learning – Reduced Order Modelling) framework that

uses a neural network to perform non-linear projections to reduced order states and efficiently predict future time steps of fluid simulations. Another avenue is multi-fidelity modelling, where coarse and fine simulations are combined to balance accuracy and cost (42, 43); for instance, a low-fidelity model might be used for most of a simulation, with high-fidelity corrections inserted only when needed.

ROM applications in stirred tanks have indeed grown rapidly in recent years, illustrating the versatility of these models. Researchers have developed reduced-order models for gas–solid mixing processes that can predict solids suspension behaviour across various tank configurations much faster than full CFD while still maintaining high accuracy (40, 46). Others have coupled POD-based ROMs with evolutionary algorithms to optimise impeller placement (e.g. finding the best impeller axial position) far more efficiently than evaluating each design with CFD (41). Surrogate models have also been employed as ROMs for complex reactors – for example, data-driven ROMs have enabled real-time prediction of key hydrodynamic properties (such as gas holdup in bioreactors and bubble column flows) under varying operating conditions (47). These diverse applications demonstrate how ROM techniques can significantly accelerate analysis and optimisation in multiphase mixing systems.

It is important to note, however, that developing a truly accurate ROM for a multiphase stirred tank is non-trivial. The underlying flow features (vortices, bubble plumes, etc.) can be highly sensitive to operating conditions, so a ROM must be trained or constructed over a sufficiently rich set of scenarios. If a new operating point lies outside the range of the training snapshots, the ROM's predictions may degrade significantly. To avoid this, the snapshot set should cover the expected range of impeller speeds, gas flow rates, and liquid properties—ensuring that the ROM basis captures all relevant flow modes.

Another challenge is that ROMs derived via Galerkin projection can suffer stability issues: a truncated model may not conserve energy and can drift unphysically over long simulation times. To stabilize reduced models, various techniques have been proposed, such as enforcing stability constraints when truncating modes, introducing calibrated linear damping terms (closures), or adding artificial eddy viscosity at the ROM level. Researchers have also explored advanced closure strategies – from empirically tuned eddy-viscosity corrections to machine-learned closure models – to improve a ROM's long-term stability and accuracy (34, 35). Finally, regularization techniques in data-driven ROMs (for example, applying L1/L2-norm penalties or using early stopping during neural network training) are often employed to prevent overfitting to the training data.

Overall, ROM technology offers a powerful complement to high-fidelity CFD. In the specific area of gas–liquid stirred tanks, ROMs enable rapid evaluation of design alternatives and real-time process control that would be impossible with direct CFD alone. Many recent works have demonstrated that with careful construction and validation, ROMs can predict key flow characteristics (like free-surface shape, circulation patterns, and mixing metrics) with acceptable accuracy, at simulation speeds that are orders of magnitude faster than CFD. The continued development of ROM techniques – including hybrid approaches that incorporate physical laws and advanced data-driven models – is an active research frontier. In the next section, it will be discussed that the remaining challenges and gaps identified in the literature, which motivate further improvements in both high-fidelity CFD and ROM approaches for multiphase mixing systems.

2.5 Remaining Challenges and Research Gaps

Despite significant advancements in CFD and ROM approaches, several challenges remain in accurately and efficiently modelling gas–liquid stirred tanks. One major challenge is scaling up from lab-scale studies to industrial-scale tanks. As tank volume increases, maintaining geometric, kinematic, and dynamic similarity becomes difficult – large tanks often cannot simultaneously preserve Reynolds, Froude, and Weber number equivalence with smaller ones (24). Traditional scale-up rules (such as keeping constant power per volume or constant impeller tip speed) frequently fail to reproduce the same flow regime at different scales (25). For example, a vortex that is well-contained in a small pilot vessel may deepen and entrain gas in a much larger vessel even if dimensionless groups are matched, due to subtle geometry effects and Reynolds number differences (24). This uncertainty means that purely relying on CFD at one scale to predict another can be risky without thorough validation. A related gap is the lack of standardized validation data for large-scale multiphase mixing. While many lab-scale experiments (PIV velocity fields, high-speed camera visualization of free surfaces, etc.) are reported in the literature, comparable data at industrial scale are scarce. This hampers the validation of CFD/ROM models intended for full-scale prediction. Collaborative efforts between academia and industry to create large-scale benchmark datasets would greatly aid model development (25).

Another challenge lies in turbulence–interface coupling in simulations (29,32). Most RANS turbulence models were developed for single-phase flows and do not inherently account for the dampening of turbulence at a deformable gas–liquid interface (32). Ad-hoc fixes (like the curvature correction or damping functions mentioned earlier) are not rigorously validated for all conditions. Large Eddy Simulations can resolve some of these interactions explicitly, but at prohibitive cost for design studies (36). Accurately

modelling how turbulence is suppressed as the liquid approaches an interface (due to lower eddy velocities in the gas and surface tension effects) is an open problem – too much turbulence leads to over-predicted surface ripple; too little gives an artificially smooth surface (32). Rodi's work suggests tuning turbulent viscosity near interfaces, but a more universal model or method for this across flow regimes is lacking (36). Similarly, multiphase turbulence closures (e.g. dispersed k–epsilon models for bubbly flows) require empirical coefficients that may not transfer well between systems. The interfacial momentum exchange (drag, etc.) and breakup/coalescence models in Euler–Euler simulations also need more fundamental research; current models are often calibrated case-by-case.

On the ROM front, a key gap is ensuring robustness and generalizability of reduced models for multiphase flows (34,35). ROMs for single-phase flows (e.g. cavity or channel flows) are relatively mature, but multiphase ROMs must contend with moving interfaces, topological changes, and possibly discontinuous solution fields (34). One issue is how to construct ROM basis functions that adequately represent both bulk flow structures and the interface dynamics (39). If one simply uses velocity and pressure snapshots, important interface information (like volume fraction) might not be captured well. Some researchers have explored including the phase indicator field in the POD state vector or using separate bases for each phase (40). The optimum approach for this is still unclear and likely system-dependent. Another issue is stability and closure of ROMs: low-order models derived from projection can be unstable, as discussed, and this is exacerbated in multiphase flows with strong nonlinearity (36). Ahmed et al. pointed out that without proper closure models or regularization, ROMs may diverge or produce unphysical results over long simulations (34). Ensuring stability often means introducing some form of modelled dissipation into the ROM equations, which if

overdone can erase the very advantages of ROM (i.e. it might damp out the dynamic modes too strongly). Data-driven ROMs also face the challenge of extrapolation – a trained model might perform poorly if the flow enters a regime not represented in the training data. Developing adaptive ROMs that can recognize when they are outside their valid domain and possibly request a full-order simulation update (a concept of multi-fidelity simulation) is an ongoing research area. Furthermore, the integration of ROMs with control systems in an online fashion requires them to be not only fast but also reliable in real time, with known error bounds. Quantifying the uncertainty in ROM predictions (e.g. via uncertainty quantification techniques or randomized training approaches) is another gap that needs to be addressed for industry adoption.

In summary, the literature survey reveals that while current CFD methods (particularly VOF-based RANS simulations) are capable of capturing the primary behaviour of gas–liquid stirred tanks, they require significant expertise to apply and still face limitations in extreme conditions (29). Euler–Euler models provide an alternative for highly aerated systems but depend on empirical closure laws that may not universally apply. Reduced-order models show great promise to alleviate computational burdens, enabling fast simulations for scale-up and control, but they are not yet "plug-and-play" for complex multiphase systems – careful tailoring and validation of ROMs for each new scenario is necessary. Key research gaps include: (1) improved turbulence modelling near gas–liquid interfaces; (2) better multiphase model calibration strategies (perhaps assisted by AI or optimisation) to ensure predictive accuracy without exhaustive experimentation; (3) standardized large-scale validation datasets for model testing; and (4) enhanced ROM frameworks that incorporate multiphase physics and guarantee stability. Addressing these gaps will likely involve cross-cutting efforts, combining high-fidelity simulations, advanced experimental measurements, and novel modelling

techniques. The continued development of CFD and ROM methodologies, as well as their hybridization, is critical for advancing the design and scale-up of stirred tank reactors.

2.6 Summary of Literature Review

In this chapter, it has been reviewed the state of the art in both high-fidelity CFD and reduced-order modelling for gas–liquid stirred tanks. Conventional CFD approaches have evolved to capture increasingly complex physics: modern interface-capturing methods like VOF can reproduce free-surface vortices and coupling with discrete phases, while advanced turbulence closures (SST, RSM) improve accuracy for swirling flows. Still, full CFD simulations are computationally intensive and can be sensitive to modelling choices, especially for large-scale or highly aerated systems. It discussed how interface modelling strategies must be chosen based on flow regime – VOF for coherent interfaces, Euler–Euler for dispersed bubbles – and how turbulence and interphase interactions remain challenging to model across all scales. This case also surveyed reduced-order models as a promising solution to CFD's limitations. By extracting dominant flow modes (e.g. via POD) and building low-dimensional surrogate models, ROMs enable rapid predictions with minimal loss of accuracy. Several studies have successfully applied ROMs to stirred tanks, achieving substantial speed-ups and demonstrating potential for real-time simulations and optimisation. However, the literature also highlights important gaps: the need for better handling of turbulence near interfaces, improved validation at industrial scales, and more robust ROM techniques tailored to multiphase flows

Section 3. METHODOLOGY

Various multiphase flow models, such as the Euler-Euler method, the discrete element method (DEM), and the volume of fluid method (VOF), have been well validated in stirred tank simulations. In this study, a high-fidelity full-scale computational model was developed to generate multiple design points, each representing a different operating scenario. Based on this comprehensive model, flow field snapshots were collected at multiple impeller speeds. These snapshots were then used to configure a reduced-order model in ANSYS, using its dedicated reduced-order modelling add-on module to load the data and build a compact model to capture the essential dynamic characteristics of the full simulation. By extracting the main flow patterns from the high-fidelity data set, the reduced-order model provides a low-dimensional representation that enables rapid evaluation of the system behaviour under different conditions. The overall workflow of snapshot generation and reduced-order model implementation is shown in Figure 1(31).

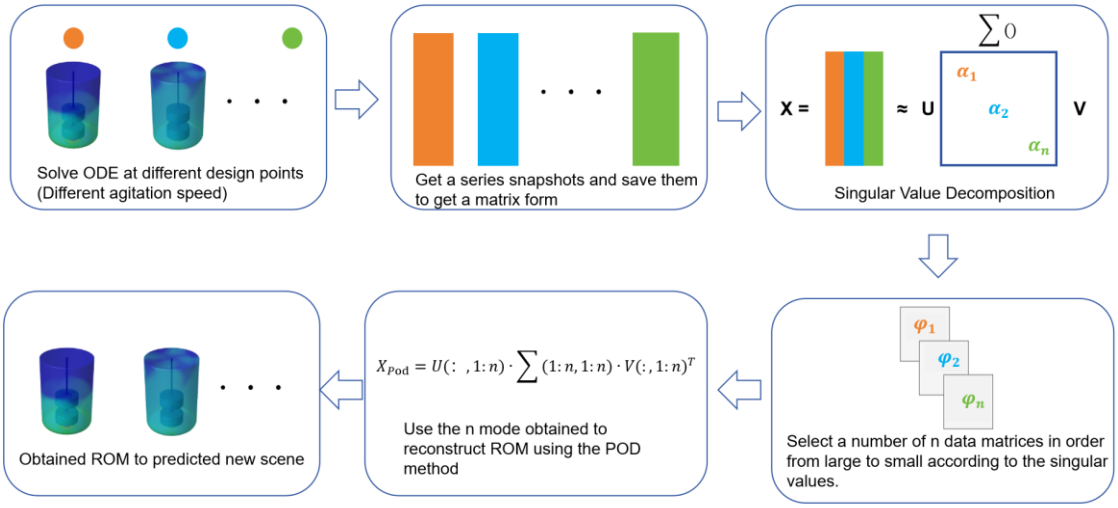

Figure 1. ROM implementation process: The full-order model is solved at multiple design points (i.e., specific combinations of input parameters, such as impeller speed and fluid properties) to generate snapshots, which are then used to construct the reduced-order model.

3.1 Implementation Steps

3.1.1 Snapshot Acquisition

In the initial stage, a series of full-order models (FOMs) are constructed to generate simulation snapshots. Each snapshot corresponds to the steady-state solution of the multiphase flow simulation at a specific rotational speed, using the volume of fluid (VOF) formulation to capture the interfaces between the phases. The governing equations are discretized using the finite volume method: the computational domain is divided into control volumes and each conservation equation (mass, momentum, etc.) is integrated over the volume to ensure local conservation of the flow. This discretization transforms the partial differential equations into a set of algebraic equations that are iteratively solved to obtain a high-fidelity solution. The result is a full-order VOF model that contains detailed information about the multiphase flow field. By repeating this process at different rotational speeds, a series of multiphase flow models can be obtained, each representing a different operating condition.

3.1.2 Singular Values Extraction

ANSYS uses its built-in functions to convert the simulation snapshot data into a matrix. In this matrix, each column corresponds to the steady-state condition of the multiphase flow at a specific rotational speed, and each row corresponds to a key flow parameter (such as mass, density, static pressure, shear force, volume fraction, etc.). The matrix is

then decomposed using singular value decomposition (SVD) to obtain singular values and corresponding singular vectors. The magnitude of each singular value reflects the relative importance or contribution of the associated feature (mode) to the entire flow system. By arranging the singular values in descending order, the most important features can be identified and prioritized. This process achieves dimensionality reduction and feature selection by retaining only the main features (modes) and discarding features associated with negligible singular values, thereby reducing data complexity while retaining the most critical information about the flow system.

3.1.3 ROM Reconstruction

During the reduced-order modelling (ROM) process, the largest n singular values and their associated modes are selected to define a reduced-order basis. ANSYS then automatically applies this POD-based basis to reconstruct the system dynamics, thereby building a ROM. The resulting model can predict new scenarios. Selecting only the largest singular values effectively reduces the dimensionality of the system because smaller singular values correspond to less important modes and are discarded. This ROM is applied to a mixing tank rotating at 240 rpm as a representative example. The selected outputs of interest—mass, density, static pressure, gas and liquid volume fractions, turbulent dissipation rate, and wall shear stress—determine how many modes are retained. In general, the number of singular values (modes) retained should be at least as large as the number of output variables to ensure that the model accurately captures all outputs.

3.1.4 Factors Influencing ROM

Creating a reduced-order model requires a series of snapshot data generated by the full-order model. Therefore, a reliable full-scale model parameter setting is required. This

setting can be used to obtain a series of snapshot data of the stirred tank at different speeds. Model selection for simulating multiphase flow. According to the requirements of the comparative experiment, the changes in the liquid level were observed and compared. A multiphase flow model that can clearly represent the changes in the liquid level is required. Many existing studies have shown that the VOF model can clearly capture the changes at the interface between the two phases.

3.2 Governing Equations for VOF and Turbulence

This section states the model equations used for a two-phase (gas–liquid) flow with an interface captured by the Volume of Fluid (VOF) method and closed with RANS turbulence models. The presentation follows a single-velocity, shared-pressure VOF formulation where material properties vary with the phase fraction. The intent is to provide a self-contained reference for the conservation laws, interface physics, turbulence closure, and typical boundary conditions used for stirred tanks or other agitated vessels, but the equations are general and can be adapted to channels, risers, or free-surface reactors.

3.2.1 Mixture quantities and phase fraction

The VOF method introduces a scalar volume fraction α (x, t) for the liquid phase. In each computational cell, α=1 denotes pure liquid, α=0 denotes pure gas, and intermediate values imply that the interface cuts the cell. Under the shared-velocity assumption, the density and dynamic viscosity entering the momentum equation are mixture quantities defined by volume averaging:

• $\rho(\alpha) = \alpha \rho_l + (1-\alpha) \rho_g$

• $u(\alpha) = \alpha u_l + (1-\alpha) u_g$

Here ρ_l, ρ_g are the constant densities of the incompressible liquid and gas; u_l, u_g are their dynamic viscosities. In many codes, the turbulent viscosity u_t is added to obtain an effective viscosity u_eff = u + u_t in the diffusive fluxes of momentum and turbulence scalars.

3.2.2 Mass conservation

For a mixture of two incompressible phases advected by a single velocity field u (x, t), mass conservation is expressed as:

$\partial\rho/\partial t + \nabla\cdot(\rho\, u) = 0$

When both phases are individually incompressible and only the composition ρ(α) varies, this reduces to the solenoidal constraint on the velocity field if the α-transport is solved consistently:

$\nabla\cdot u = 0$

3.2.3 Momentum conservation

The mixture momentum equation (in an inertial frame) is:

$\partial(\rho\, u)/\partial t + \nabla\cdot(\rho\, u \otimes u) = -\nabla p + \nabla\cdot[\, u_{eff}\,(\nabla u + \nabla u^{T})\,] + \rho\, g + F_{\sigma} + S_{body}$

Here p is the mechanical pressure, g is gravity, and F_σ is the continuum surface force (CSF) representing surface tension at the gas–liquid interface. S_body may include optional source terms such as rotating-frame effects in a multiple reference frame (MRF) region around an impeller. In a rotating frame with angular velocity Ω, the additional terms are $-\rho\,\Omega\times(\Omega\times r) - 2\rho\,\Omega\times u$.

The effective viscosity u_eff = u + u_t includes the turbulent viscosity u_t supplied by the turbulence model. The production of turbulent kinetic energy is linked to the symmetric strain-rate tensor $S = (\nabla u + \nabla u^{T})/2$.

3.2.4 Interface capturing and surface tension (CSF)

The phase fraction obeys a conservative transport equation with an optional compression term to keep the interface sharp:

$\partial\alpha/\partial t + \nabla\cdot(\alpha\, u) + \nabla\cdot[\, \alpha\,(1-\alpha)\, u_c\,] = 0$

Here u_c is an artificial compression velocity aligned with the interface normal $n = \nabla\alpha / |\nabla\alpha|$ and confined to cells with $0<\alpha<1$. The last term counteracts numerical diffusion of the interface while preserving boundedness.

Surface tension is modelled by the CSF approach as a volumetric force concentrated near the interface:

$F_\sigma = \sigma\, \kappa\, \nabla\alpha$

where σ is the (possibly temperature- or composition-dependent) surface tension and κ is the interface curvature, $\kappa = -\nabla\cdot n = -\nabla\cdot(\nabla\alpha / |\nabla\alpha|)$. For wall-bounded flows with a specified static contact angle θ, the interface normal near the wall is corrected so that $n\cdot t$ satisfies the geometric condition implied by θ, which influences capillary rise and meniscus shape.

3.2.5 Turbulence closure (RANS)

Reynolds-averaged Navier–Stokes (RANS) models provide u_t to represent the momentum transport by unresolved eddies. Two widely used closures for separated, swirling, or free-surface flows are the k–e family and the k–ω SST model. Below, ρ and u refer to mixture properties; model constants take their conventional values unless otherwise stated.

3.2.6 Standard k–e model

The standard k–epsilon model solves transport equations for the turbulent kinetic energy k and its dissipation rate e:

$\partial(\rho k)/\partial t + \nabla\cdot(\rho u k) = \nabla\cdot[(u + u_t/\sigma_k) \nabla k] + P_k - \rho e + G_b$

$\partial(\rho e)/\partial t + \nabla\cdot(\rho u e) = \nabla\cdot[(u + u_t/\sigma_e) \nabla e] + C_{1\varepsilon} (e/k) P_k - C_{2\varepsilon} \rho e^2 / k + C_{3\varepsilon} (e/k) G_b$

with $u_t = C_u \rho k^2 / e$. The production term is $P_k = 2 u_t S:S - (2/3) \rho k \nabla\cdot u$ (the last term vanishes for incompressible flow). The optional buoyancy production G_b accounts for density gradients aligned with gravity (important in stratified gas–liquid regions). In the present simulations, standard wall functions are employed for near-wall treatment. These functions are applicable for y^+ values in the range of 30–300. The mesh used in this study was generated to maintain y^+ within this range, as detailed in Section 4.2.

3.2.7 Realizable/ RNG variants of k–e

Realizable and RNG k–e models modify u_t and the e-equation coefficients to improve performance in swirling and separating flows. The realizable form enforces mathematical constraints on the Reynolds stresses, while RNG adds scale-dependent terms derived from renormalization. If strong swirl from an impeller is present, these variants often reduce overprediction of eddy viscosity compared to the standard model.

3.2.8 k–ω SST model

The k–ω SST model blends the near-wall accuracy of k–ω with the free-stream robustness of k–e. It solves for k and the specific dissipation rate ω, includes a cross-diffusion term, and limits u_t to avoid unphysical shear-layer growth:

$\partial(\rho k)/\partial t + \nabla\cdot(\rho u k) = \nabla\cdot[(u + \sigma_k u_t) \nabla k] + P_k - \beta^* \rho k \omega$

$\partial(\rho\,\omega)/\partial t + \nabla\cdot(\rho\,u\,\omega) = \nabla\cdot[\,(u + \sigma_\omega\,u_t)\,\nabla\omega\,] + \alpha\,(\omega/k)\,P_k - \beta\,\rho\,\omega^2 + 2(1-F_1)\,\rho\,\sigma_{\omega 2}\,(1/\omega)\,\nabla k\cdot\nabla\omega$

$u_t = \rho\,a_1\,k / \max(a_1\,\omega, S\,F_2)$

F_1 and F_2 are blending functions switching between k–ω and k–e behaviour across the boundary layer; α, β, β^*, σ_k, σ_ω, $\sigma_{\omega 2}$, and a_1 are model constants. SST is commonly preferred for free-surface flows with strong adverse pressure gradients.

3.2.9 Reynolds stress transport (RSM)

When anisotropy is crucial (e.g., strong streamline curvature behind baffles or in the impeller discharge), a Reynolds Stress Model can be used. RSM advances transport equations for the six independent components of the Reynolds stress tensor R_{ij}, coupled to a scale-determining equation (e.g., e or ω). While more expensive, RSM can better capture secondary circulations and normal-stress imbalances that drive free-surface deformation.

3.2.10 Pressure–velocity coupling and time integration

In segregated solvers, pressure–velocity coupling is achieved by SIMPLE, SIMPLEC, or PISO. The VOF transport is advanced alongside the momentum equations with bounded high-resolution convection schemes to maintain $0 \le \alpha \le 1$. For transient problems, the time step is constrained by both a global Courant number ($Co = \Delta t\,|u| / \Delta x$) and a capillary time-step limit ($\Delta t < \sqrt{\rho\,\Delta x^3 / (2\pi\,\sigma)}$) to resolve capillary waves and avoid interface ringing. Under-relaxation and multigrid can assist convergence for stiff surface-tension-dominated flows.

3.2.11 Boundary and initial conditions

Typical settings for a stirred tank or open reactor are: (i) no-slip walls for velocity; (ii) either high-Re wall functions or low-Re treatment depending on y+; (iii) pressure outlet

or open-to-atmosphere boundary at the free surface with α specified (often α=1 for liquid below and α=0 for gas above at t=0) and a static contact angle θ on solid walls; (iv) at gas spargers or inlets, specify mass or volumetric flow and α=0; (v) initial fields for u, k, e/ω consistent with the quiescent or ramped start-up condition. In MRF, rotating-zone interfaces use conservative fluxes; for sliding mesh, grid interfaces exchange fluxes across the nonconformal boundary each time step.

3.2.12 Dimensionless groups and scaling

To interpret results and support scale-up, it is useful to monitor: Reynolds number Re = ρ N D^2 / u (impeller speed N and diameter D), Froude number Fr = N^2 D / g (importance of inertia vs gravity for vortex depth), and Weber number We = ρ N^2 D^3 / σ (balance of inertia and surface tension). In gas–liquid systems, the gas flow number Fl_g = Q_g / (N D^3) and power number N_P = P / (ρ N^3 D^5) also help organize data.

3.2.13 Assumptions and applicability

The above VOF–RANS framework assumes Newtonian phases, negligible phase change, and a resolved but diffuse interface over a few cells. At very high gas volume fractions (e.g., churn–turbulent), an Euler–Euler two-fluid model may be more appropriate. For flows where large-scale unsteadiness dominates (e.g., periodic blade passage), unsteady RANS or hybrid RANS–LES can reduce model-form error. The turbulence model choice should be justified a priori with reference data and a posteriori by verification/validation and sensitivity analyses.

3.3 Proper Orthogonal Decomposition Method

Currently, a widely adopted dimensional reduction technique when creating reduced-order models is proper orthogonal decomposition (POD). Its purpose is to reduce the reconstruction error produced when a series of snapshots is projected onto a set of initial functions, thereby it can reduce the order of system (50). To achieve this minimization, the left singular vector is obtained from using SVD to decompose the snapshot matrix form. The equation can be showed as below

$$X = U\Sigma V^{T} \tag{14}$$

The matrix on the right contains the matrix of left vectors on the left, the transposed matrix of the left vector on the right, and the singular values on its diagonal in the diagonal matrix. And the matrix reconstructed by the POD is a good fit for the stirred tank of this study, many singular value is close to value 0, and this is the source of the corresponding columns that can be discarded. The POD' s equation which choose to be kept are saved in the form that needs to be represented. At the same time, the POD coefficient used for reconstruction can be obtained by remultiplying the data obtained from the snapshot:

$$(u^{recon})^{k} = RR^{T}u^{k} \tag{15}$$

And u^{k} is the k th matrix of snapshot and $(u^{recon})^{k}$ is set as a reconstruction of them. So, the reconstruction error from a series of N snapshots matrix form $\{u^{1}, u^{2}, u^{3}, \ldots, u^{N}\}$ can be showed as

$$\frac{1}{N}\sum_{k=1}^{N}(u^{k} - (u^{recon})^{k}) \cdot (u^{k} - (u^{recon})^{k}) \tag{16}$$

Typically, in previous steps of ROM creation, before using the SVD method, the average value is subtracted from the snapshot. In this research, it was found that it does not have too much effect. In this case, the multi-phase flow is stirred in a stirred tank and the snapshot matrix contains three velocity components,

$$u^k = (u_1^k, u_2^k, u_3^k, \ldots, u_M^k, v_1^k, v_2^k, \ldots, v_M^k, w_1^k, w_2^k, \ldots, w_M^k, \alpha_1^k, \alpha_2^k, \ldots, \alpha_M^k) \quad (17)$$

And w_i^k and α_i^k are the agitation speed and the liquid height at the Mth node of the kth matrix of snapshot. In this case, the velocity components are scaled to the range set by the input [-1,1] so that the magnitude of velocity is similar to the magnitude of the volume fraction in each grid.

Section 4. Data source and Experimental Cases detail

4.1 Overview of Data Source

This study utilises both experimental and simulation data for model development and validation. The experimental data were obtained by digitizing a published plot (Figure 9) from Ciofalo et al. (1996) using Plot Digitizer software. In this process, the axes of the original figure were calibrated with known scale values, and points on the plotted curves were selected to reconstruct the numeric dataset (24). This provided a series of measured free-surface elevations as a function of radial position for an unbaffled stirred tank.

Corresponding data from CFD simulations (ANSYS Fluent 2021) were generated for comparison. The liquid free-surface profile (i.e. the liquid height versus radius) was extracted directly from the simulation by defining an appropriate coordinate system: the horizontal axis was set as radial distance from the tank centreline, and the vertical axis as the fluid surface height at that radius. Using these coordinates, the VOF-based simulation output yielded the water surface elevation as a function of radial position. This procedure ensured that the simulated free-surface shape could be directly compared to the experimental profile.

4.2 Rushton Tank Geometry Details

To validate the ROM predictions, a high-fidelity full-order CFD model was first constructed and run for the gas–liquid system in an unbaffled Rushton tank. This Volume of Fluid (VOF) simulation captured the key flow phenomena, recording the transient deformation of the liquid free surface (vortex shape) and the static pressure distribution in the tank under impeller agitation. The geometric parameters of the tank

and impeller were based on the comprehensive data set provided by Haque et al. (2006) (2). The vessel is a flat-bottom cylindrical tank of diameter T = 0.19 m with an initial liquid fill height H = T (i.e. 0.19 m). It is equipped with a standard six-blade Rushton turbine impeller of diameter D = T/2 (i.e. 0.095 m). The impeller is positioned at a clearance C = T/3 (approximately 0.063 m) from the tank bottom (measured to the impeller disk's mid-plane). Table 1 summarises the detailed tank dimensions.

The computational domain was discretized using a hybrid unstructured mesh generated by ANSYS Meshing 2021 R1. The main mesh region consisted primarily of hexahedral elements, with prism layers applied along the tank wall and impeller surface to resolve the boundary layer gradient. Local mesh refinement was employed in high-velocity gradient regions:

- Impeller outlet region,
- Free surface region (expected gas-liquid interface),
- Near the baffle.

The final mesh contains approximately 4.36×10^5 elements, with a maximum skewness below 0.85 and orthogonality above 0.15, meeting the recommended stability and accuracy standards for multiphase RANS simulations (61). Near-wall resolution was controlled within the range of 30–300, consistent with the use of standard wall functions. To ensure the solution is independent of spatial discretization, a mesh independence study was conducted in Chapter 6.2, using three successively refined meshes: coarse (approximately 2.5×10^5 elements), medium (approximately 4.36×10^5 elements), and fine (approximately 7.2×10^5 elements). The predicted free surface profile and velocity field showed negligible differences (less than 2%) between the medium and fine meshes,

confirming that the medium mesh provided sufficient resolution. Therefore, all subsequent simulations used the medium mesh. This method conforms to best practices for stirred tank CFD reported in reference (67). Experimental results from the literature were used to ensure the CFD model's credibility. In particular, the free-surface shape and other measured quantities reported by Haque et al. (2006) were employed for comparative validation of the simulation (7). By using the same geometry and operating conditions as Haque's study, the CFD model's predictions of liquid surface deformation and other flow characteristics could be directly compared against experimental observations, lending confidence that the model accurately reproduces the real system's behaviour.

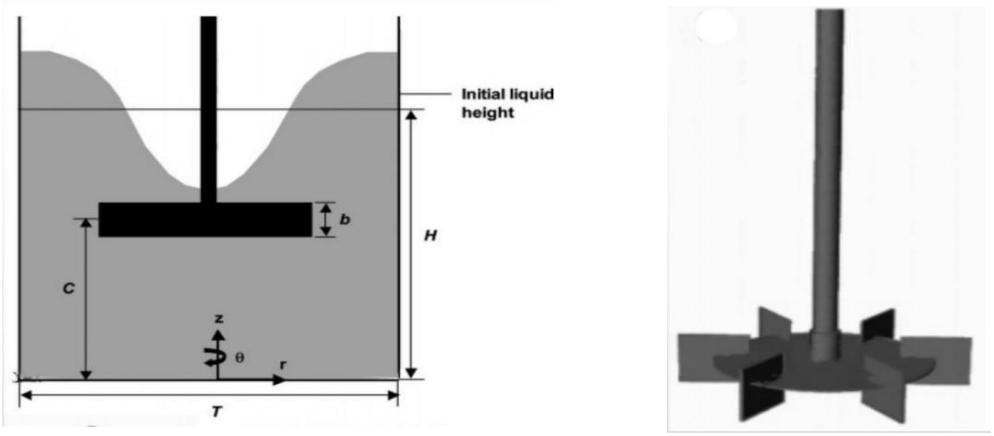

Figure 2. Configuration of an unbaffled agitated vessel with a six-bladed Rushton turbine

vessel geometry	Rushton turbine
vessel diameter (T)	0.19m
impeller diameter (D)	0.095m(T/2)
impeller clearance (C)	0.063m(T/3)
blade height (b)	0.019m
blade width (a)	0.023m
initial liquid height (H)	0.19m

Figure 3. Detail of an unbaffled agitated vessel with a six-bladed Rushton turbine

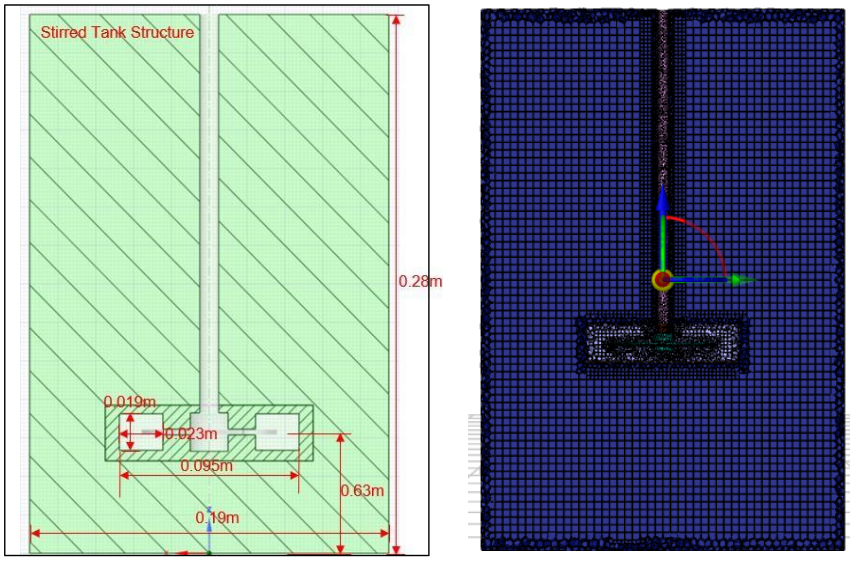

Figure 4. Geometry detail and computational mesh of the stirred tank

The above figure 3 shows the dimension detail, and the created model is created based on the dimensions in the above figure.

4.3 Material Specifications

All fluid properties in the simulation were selected to represent standard air–water systems at ambient conditions, consistent with values reported in the literature (49, 61). The gas phase was air at 20 °C and 1 atm, with a density of 1.225 kg/m³ and a dynamic viscosity of 1.789×10^{-5} Pa·s. The liquid phase was water at 20 °C, with a density of 998.2 kg/m³ and a dynamic viscosity of 1.003×10^{-3} Pa·s. The surface tension at the air–water interface was set to 0.072 N/m, which is the standard value for clean water at 20 °C (61). These properties are consistent with those used in previous experimental and numerical studies of unbaffled stirred tanks (7, 49).

	Material	Density	Dynamic viscosity	Surface tension
Gas phase	Air	1.125 kg/m³	1.78e-5 kg/m³ s	72.8 m N/m
Liquid phase	Water	997 kg/m³	1.0013 kg/m³ s	72.8 m N/m

Figure 5. Material properties

4.4 Agitation Speed Details

The direction of rotation of the impeller is clockwise. And the rotation speed (N) of the control group: 139, 194 and 240 rpm;

And the Reynolds number calculation formula:

$$Re = (\rho * N * D2) / u \quad (19)$$

where:

ρ is the density of water (assumed to be 1000 kg/m^3)

N is the rotational speed of the impeller in revolutions per minute (rpm)

D is the diameter of the tank (assumed to be 0.19 m)

u is the dynamic viscosity of water (assumed to be 0.001 Pa*s)

Through the Reynolds number calculation formula, the calculated Reynolds numbers of the three rotational speeds are: 283819, 394706, 487059. The free surface profile of the liquid is measured at three rotation speeds.

4.5 Application of the MRF Method for Geometry Creation

The Multiple Reference Frame (MRF) approach was employed in the CFD setup to model the impeller's rotation in a steady framework. In this approach, the computational domain is divided into two zones: a rotating impeller region and a stationary outer region (the tank). The impeller zone (containing the Rushton turbine) is assigned a constant angular velocity corresponding to the impeller speed, while the outer zone remains stationary. At the interface between these two zones, momentum exchange is handled via frame transformation without physically moving the mesh. Using the MRF method thus allows the influence of the impeller rotation to be included in the flow solution (through additional source terms in the momentum equations) without the overhead of a transient sliding mesh simulation. This simplification is standard for steady-state simulations of stirred tanks (4) and provides an efficient yet accurate way to capture the time-averaged effect of impeller motion on the flow. In this study, all VOF simulations adopted the MRF scheme to model impeller-induced flow in the tank.

4.6 ROM Input and Output Description

The reduced-order model (ROM) constructed in this work is parameterized by key input variables reflecting the operating conditions. Specifically, the ROM takes as input the impeller rotational speed, the initial liquid level (fill height) in the tank, and fluid properties such as liquid density and liquid viscosity (which could vary for different

fluid scenarios, though in this work cases these were fixed to water at a given temperature). By adjusting these inputs, the ROM can be queried for different what-if scenarios within the range of conditions it was trained on. The outputs of the ROM include the primary flow field characteristics of the gas–liquid system: for example, the volume fraction distribution (gas vs. liquid) throughout the domain, the velocity field in the tank, and the shear stress distribution in the fluid. The ROM also captures derived quantities such as the effective local density of the two-phase mixture (which varies with the local gas volume fraction). These outputs provide essential information on the multiphase flow dynamics – for instance, the predicted free-surface shape, flow circulation patterns, and stress levels – which can be analysed to gain insight into the system's behaviour under different conditions. In summary, the ROM serves as a fast surrogate that, given key input parameters, predicts the flow outcomes (liquid height profile, pressure distribution, etc.) that would otherwise require a full CFD simulation to obtain.

Section 5. Simulation result analysis

This study uses quantitative error metrics to evaluate the performance of reduced-order models (ROMs) compared to experimental and full CFD results and investigates how key parameters affect accuracy. ROM predictions consider two complementary error metrics: maximum error and mean error. Maximum error is defined as the maximum absolute difference between the ROM prediction and the experimental value along the target profile (e.g., free surface height versus radial position). Mean error is the average of the absolute differences at evenly distributed points along the profile (in this case, eight points from the tank centre to the wall). Together, these metrics capture both local extreme deviations and global consistency, adhering to best practices for fluid flow ROM validation (35). A small maximum error is crucial for capturing key local features (e.g., the deepest point of a vortex), while a low mean error indicates that the ROM accurately reproduces the global free surface trend (35).

Using these two metrics allows for a robust assessment of ROM prediction performance and helps guide model improvement strategies (22). The largest ROM differences occur near the vortex centre (where flow curvature and turbulence effects are strongest), with smaller differences near the tank wall. In the intermediate region, the ROM predictions closely follow the experimental curves. This pattern is consistent with previous observations that near-interface physics and swirling flows are difficult to capture in reduced-order models (35). At operating points where the ROM's training snapshots adequately cover the flow regime, the average error remains low, indicating good global fidelity. However, isolated large errors indicate localized phenomena that neither the ROM nor the underlying CFD snapshot data can fully resolve (59). These findings

emphasise the need to improve full-order model fidelity in critical regions and ensure sufficient training data, rather than simply increasing the complexity of the ROM.

5.1 Analysis of Liquid Level Results

A high-fidelity CFD simulation was performed using the Volume of Fluid (VOF) multiphase model to resolve the gas–liquid interface in an unbaffled Rushton tank. The tank geometry, mesh resolution, and impeller rotation scheme (Multiple Reference Frame, MRF) were identical to those described in Chapter 4, reflecting a standard lab-scale configuration (51) and employing a steady-state impeller modelling approach that is widely used for stirred tanks (52). Air and water were used as the gas and liquid phases at ambient conditions, with gravity included to allow the free surface to form a steady vortex under rotation. Turbulence was modelled with the standard k–epsilon two-equation RANS model (with standard wall functions), appropriate for the high Reynolds number regime. Overall, this CFD setup provided a detailed full-order baseline solution for the multiphase flow, capturing the liquid circulation patterns and free-surface deformation due to impeller action.

Using the CFD model as a source of training data, a Proper Orthogonal Decomposition (POD)-based reduced-order model (ROM) was constructed in a commercial simulation environment. The ROM was trained on steady-state solution snapshots from the CFD simulations covering multiple operating conditions of interest. After construction and validation, the ROM could predict flow fields and free-surface shapes for new conditions significantly faster than performing a full CFD simulation.

The ROM's predictive capability was evaluated at three impeller speeds (139, 194, and 240 rpm), representing low, medium, and high agitation rates. These particular speeds coincide with those used in earlier investigations on unbaffled stirred tanks (53)(54).

All three conditions are in the fully turbulent regime (Re ≈ 2.8×10^5 at 139 rpm, 3.9×10^5 at 194 rpm, and 4.9×10^5 at 240 rpm, based on tank diameter), and each produced a pronounced vortex at the gas–liquid interface. Increasing the impeller speed from 139 to 240 rpm deepened the central vortex and altered the overall liquid height profile. Higher rotation speeds imparted stronger centrifugal forces, resulting in a deeper central vortex (i.e. a lower liquid level at the centre) with greater air entrainment, whereas lower speeds yielded a shallower surface depression and weaker turbulence (53). The ROM's predicted free-surface profiles for these cases were then compared against experimental data from the literature and full-order CFD results to assess the model's accuracy.

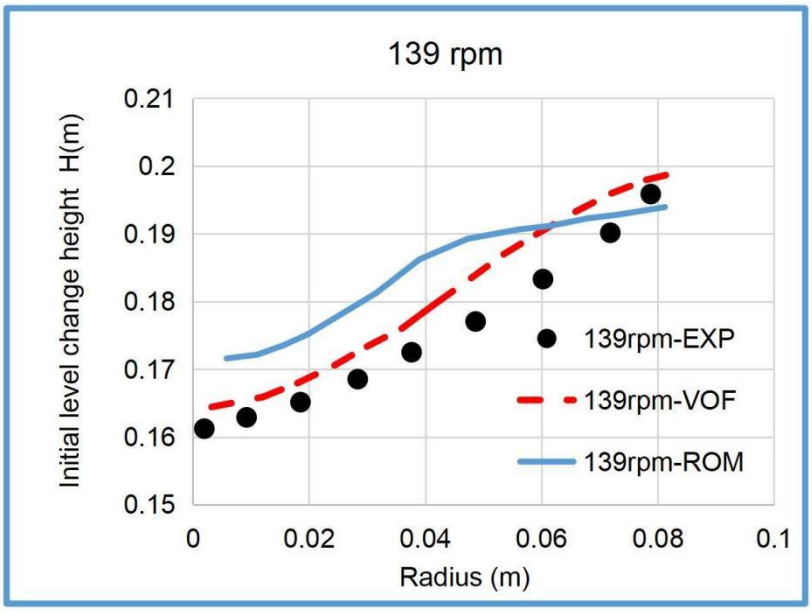

Chart 1. Comparison of Initial level change height change at 139rpm

As showed in chart 1, the free surface profile predicted by the ROM follows the correct qualitative trend, but there are some quantitative deviations from the experimental data at a speed of 139 rpm (51). The largest deviation occurs at the centre of the tank (r = 0),

where the ROM predicts a level about 0.01 m below the measured value (the central vortex is about 10 mm deeper than the observed one). This 10 mm difference is the maximum error at this low speed. Moving radially outward, the difference decreases: at the intermediate radius (~0.05 m), the error is about 0.005 m (5 mm), and near the tank wall ($r \approx 0.095$ m), the ROM and experimental water levels almost coincide (within ~0.002 m or 1-2 mm). Therefore, the ROM error is largest at the vortex core and gradually decreases towards the periphery.

In addition, the ROM free surface curve is slightly flat in the central region, while the experimental curve rises sharply with increasing radius. Nevertheless, both the ROM and the experiment show the same overall concave free surface shape, with the lowest liquid level at the centre (52). The ROM underestimates the surface curvature near the centre, but outside the core region, its predicted water levels agree well with the experimental measurements and the full CFD simulation results. Despite this small difference in low speeds, the ROM still captures the key phenomenon—the central vortex deepens at the centre and rises toward the wall—which is consistent with the behaviour of unbaffled vortex flow (52).

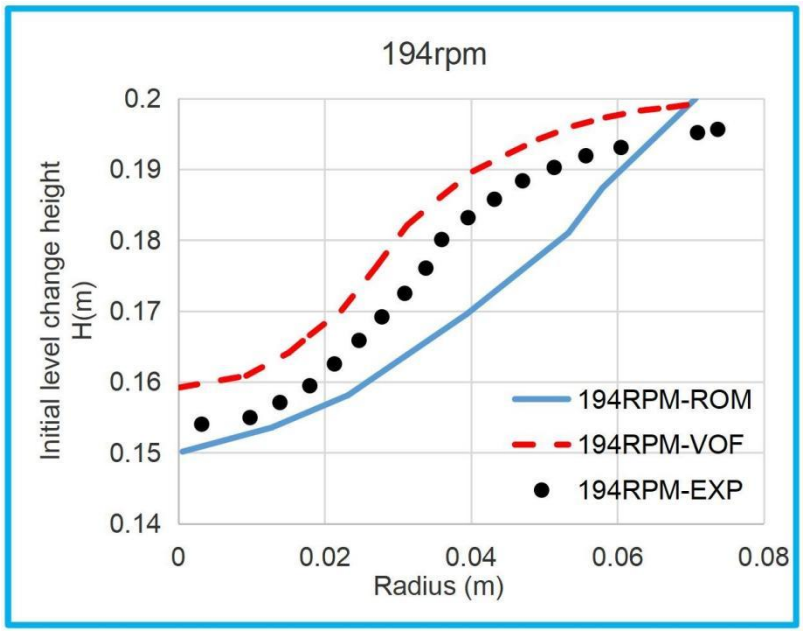

Chart 2. Comparison of the Initial level change height change at 194rpm

At 194 rpm, the ROM's predicted free-surface profile closely follows the experimental measurements overall, with only minor deviations. It matches the liquid height at the tank centre to within a few millimetres (~0.002 m error) and remains very accurate near the wall (within about 1 – 2 mm of the measured level) (53). However, at an intermediate radial position (around r ≈ 0.04 m), the ROM underestimates the free-surface height by approximately 0.007 m (7 mm) compared to the experiment – the largest discrepancy observed at this speed (53). This error indicates that the ROM predicts a slightly flatter curvature in the mid-radius region than what the experiment shows, causing the ROM and experimental profiles to diverge around r = 0.04 m (where the actual free surface has more pronounced curvature). In contrast, a high-fidelity VOF CFD simulation at 194 rpm agrees almost exactly with the experimental free-surface shape, with the CFD and measured curves overlapping within ~1–2 mm across all radii (54). In summary, at this medium rotation speed, the ROM accurately reproduces the free-surface height at the centre and near the walls, but it misses a slight bulge in the

mid-span region. The result is a localized error on the order of 6–8 mm in that mid-radius zone, highlighting it as an area where the ROM's representation could be improved (54).

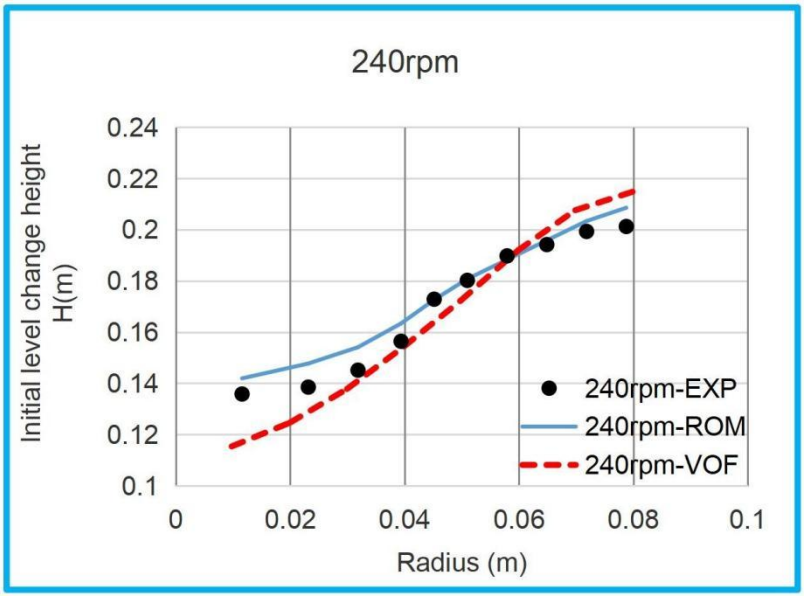

Chart 3. Comparison of Initial level change height change at 240rpm

High-Speed Case (240 rpm): according to the chart 3, the ROM's predicted free-surface profile is virtually coincident with the experimental measurements across the entire radius (51). The maximum deviation is only about 0.003 m (3 mm) at the tank centre, and elsewhere the error drops to approximately 0.002 m or essentially zero near the wall. These discrepancies are far smaller than those at 139 rpm and 194 rpm, where errors on the order of ~0.01 m (10 mm at the centre) and ~0.007 m (7 mm at mid-radius) were observed, respectively. In the 240-rpm case, by contrast, no such large deviations occur – the ROM prediction lies almost on top of both the experimental and full CFD profiles, indicating excellent agreement.

This superior accuracy at high rotation speed suggests that the ROM is capturing the flow physics more effectively under conditions of a deeper vortex and more intense

turbulence. Notably, the central vortex becomes deeper as impeller speed increases, and since the Froude number is proportional to the square of the tip speed divided by gravity, a higher impeller speed directly implies a higher Froude number. The turbulence intensity grows with the greater energy input (55). One likely reason for the improved ROM performance is that the underlying full-order CFD simulation is more reliable at 240 rpm. Stronger agitation produces a more stable, well-defined vortex that is easier for the CFD to resolve, so any mesh or turbulence-model limitations have a smaller relative impact on accuracy (7)(24). In other words, the ROM at these high-speed benefits from having high-fidelity snapshot data, resulting in a very faithful surrogate prediction.

The contrast between the high-speed and lower-speed results implies that the modest ROM errors seen at 139–194 rpm stem primarily from the training data fidelity rather than deficiencies in the ROM methodology itself. Past studies have similarly found that enhancing the CFD model (for instance, using finer grids or more advanced turbulence models) markedly improves predictions of strongly swirling free-surface flows (6). By 240 rpm, the combination of a quality CFD baseline and a sufficiently rich ROM basis captures the free-surface behaviour nearly perfectly. This demonstrates that when the full-order data are accurate, the ROM can reproduce the gas–liquid interface shape with only a few millimetres of error – an excellent level of accuracy for engineering purposes.

5.2 Design point analysis

Two of the most influential factors in building a ROM are (i) the number and choice of design points (distinct operating conditions or snapshots used for training) and (ii) the number of POD modes retained in the reduced-order basis. In the ROM construction

process (illustrated in Figure 1), both factors can be controlled: the number of design point snapshots and the number of retained modes. Intuitively, providing more design points should expand the ROM's knowledge of the system's behaviour, and including more modes should allow the ROM to capture more detail from the snapshots. In this subsection, it has been analysed how varying the number of design points affects the ROM's accuracy, using the free-surface height prediction error as the metric.

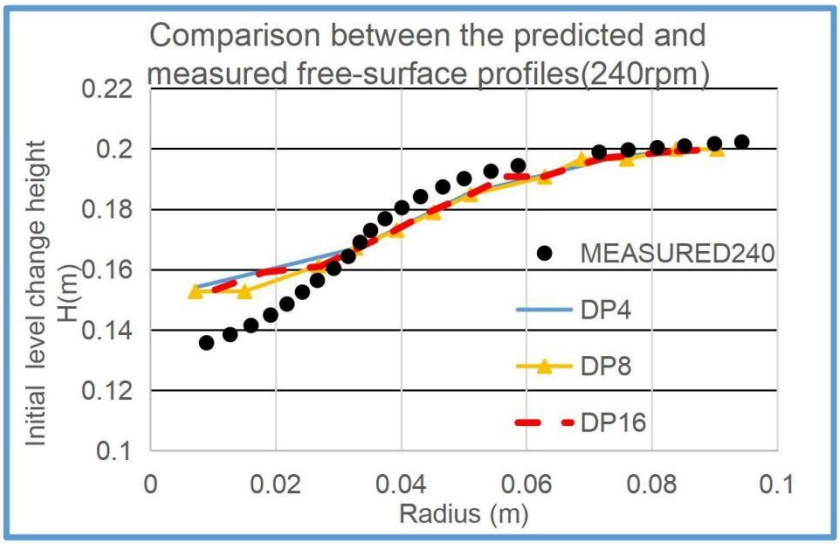

Chart 4. Comparison between the predicted and measured free-surface profiles(240rpm)

Increasing the number of design point snapshots markedly improves ROM accuracy. In this case, expanding the training set from 4 up to 8 distinct impeller speeds produced a substantial reduction in the ROM's free-surface prediction error, confirming that a broader sampling of operating conditions yields a more accurate reduced-order model (35,56). Each additional design point provides new flow information that enriches the POD basis and allows the ROM to capture more of the system's behaviour, thereby

reducing interpolation error (35). However, beyond roughly eight design points, the benefits of adding more snapshots diminish. Increasing from 8 to 12 design points yielded only a minor further decrease in error (for example, the maximum free-surface deviation dropped only slightly, from about 0.003 m at 8 points to ~0.0025 m at 12 points), and even using 16 points showed virtually no improvement (~0.002 m error) (57). In practical terms, about eight well-chosen design points were sufficient to span the dominant flow regimes in this Rushton tank case; adding more snapshots beyond this provided negligible gains in predictive fidelity (8). Overall, the design point analysis confirms that the coverage of the training set is crucial: a well-chosen moderate number of design points can yield a highly accurate ROM, whereas too few points lead to large interpolation errors, and too many points give little additional benefit.

5.3 Analysis of Mode Numbers

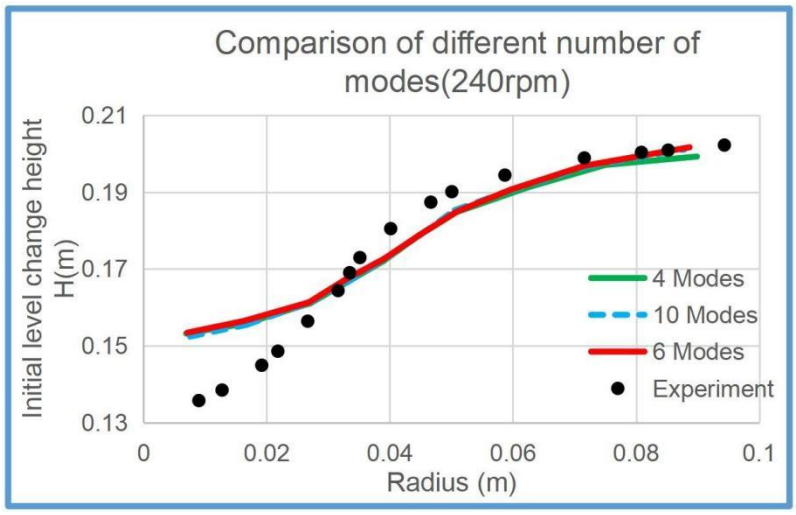

Chart 5. Comparison of different number of modes

As shown by Chart 5, increasing the number of POD modes beyond a modest count yields only marginal improvements in ROM prediction accuracy. In fact, a reduced-order model using only about 4–6 POD modes are essentially as accurate as one using

10 modes, indicating negligible gains beyond the first few dominant modes. This aligns with observations in the literature that retaining only the leading POD modes often achieves nearly the same fidelity as using many more (59). The fundamental reason is the rapid decay of POD singular values: the first few modes capture the vast majority of the flow's energy (or variance), while higher-order modes carry only a minor residual portion (35). Once these dominant energy-containing modes are included, the ROM's predictive power effectively saturates, and additional higher-order modes contribute little new information (35). In other words, the solution has essentially reached its maximum fidelity after incorporating the key modes. Furthermore, truncating the POD basis at a small number of modes confers practical advantages. A ROM built on fewer modes has a much lower-dimensional state, which significantly reduces computational cost for simulation. At the same time, using a moderate number of modes helps avoid fitting noise or insignificant flow fluctuations present in the training data (60). The highest-index POD modes (those associated with very small singular values) often represent experimental or numerical noise and minor flow structures; excluding them improves the model's generalization to new conditions by preventing overfitting to spurious details. In summary, a properly truncated POD basis not only maintains near-optimal accuracy but also yields a more efficient and robust model by avoiding the inclusion of superfluous modes.

These findings have clear implications for ROM design. They suggest that one should include only as many POD modes as needed to capture the important flow energy, and not exceed this threshold of diminishing returns. Beyond a certain cutoff, additional modes introduce unnecessary complexity without providing meaningful accuracy improvement. Thus, the recommended strategy is to select a compact set of dominant modes that collectively represent a high percentage of the system's total energy (for

example, ~95–99% of the energy content). Adopting this criterion produces a low-dimensional model that is computationally fast and stable, while avoiding the pitfalls of overfitting or wasted modes that yield minimal benefit. In essence, focusing on the quality of the modes rather than the quantity leads to a ROM that can predict system behaviour with high fidelity. A carefully chosen small basis thereby achieves the minimum effective model dimension needed for optimal performance, confirming that more is not always better in POD-based ROM design.

Section 6. ROM improvement

Although POD-based ROM has developed a huge acceleration for the complete CFD for gas stir tanks, it still shows limitations that require further improvement. In particular, simply increasing the number of POD patterns or adding more snapshot data will yield more than a certain point of return. This suggests that the rest of the errors are due to fundamental issues – mainly the fidelity of high-fidelity CFD data and the representation of training snapshots – rather than just the undersized ROM scale. To address these issues, this chapter focuses on experimental strategies to improve ROM performance against three key aspects:

1. Refining the full-order CFD model: High-fidelity CFD settings are enhanced (e.g., through finer mesh resolution and improved physical models) to generate more accurate and realistic simulation snapshots for ROM training. A more reliable CFD baseline reduces error propagation during the training phase.

2. Optimise ROM **construction**: Using the improved CFD dataset, the ROM has been rebuilt with wiser snapshot selection and an appropriate number of POD modes. By covering a wider range of operating conditions (especially when the original ROM shows big errors), ROMs can achieve higher accuracy without unnecessary complexity.

3. Incorporating data-driven enhancements: Advanced data-driven techniques, such as machine learning, are introduced to extend the predictive capabilities of ROM beyond the standard POD-Galerkin framework. For example, AI-based methods can help identify key regions in the parameter space and propose additional snapshot simulations to enrich the training set in error-prone areas, avoiding brute-force sampling. However, ensuring the generalizability of such data-driven ROMs remains an open challenge, and

hybrid approaches combining data-driven techniques with physical insights may offer a promising path forward (34, 36).

In the following sections, each of these improvements will be implemented and analysed in detail. By comparing the performance of ROM before and after each enhancement, it has been quantified the extent to which each strategy reduces prediction errors and improves stability. Overall, the enhanced ROM is significantly closer to the high-fidelity CFD (and experimental) results, while retaining the fast calculations that make the ROM so valuable.

6.1 Strategies for Error Reduction in ROM

The accuracy of a Volume of Fluid (VOF) simulation for gas–liquid flow is governed by several factors commonly highlighted in CFD best-practice guidelines. Key considerations include:

- Grid Resolution: The mesh must be fine enough to resolve important flow features and the gas–liquid interface. An overly coarse grid will smear the free surface and miss small vortices or waves. Refining the mesh (especially near the impeller and free surface) sharpens interface predictions and improves velocity gradients. Simulations should be checked for grid independence by comparing results on successively finer meshes (11). Adequate grid resolution reduces numerical diffusion and yields solutions closer to true continuum behaviour.
- Turbulence Model: The choice of turbulence closure has a major impact on the predicted flow structure and free-surface shape. Different turbulence models (e.g., variants of k–e, k–ω, or Reynolds Stress Models) can produce markedly different vortex depths, turbulence levels, and phase distributions (6)). Li et al. (9) demonstrated that in a magnetically stirred unbaffled tank, the RSM and SST k–ω models predicted free-surface shapes and velocity distributions more accurately than the standard k–ε model,

underscoring the sensitivity of free-surface predictions to turbulence model selection.). Likewise, Lane et al. (2005) reported that CFD results deviated from experiments largely due to the turbulence model's inability to capture the intense vortex properly. Selecting an appropriate turbulence model (and tuning its parameters) is therefore essential for accurate multiphase CFD results.

- Time Step Size (Transient Simulation): In unsteady VOF simulations, the time step Δt must be small enough to satisfy the Courant–Friedrichs–Lewy (CFL) condition (a stability requirement stating that during one time step the fluid motion should not move information across more than one computational cell, otherwise numerical errors or instabilities can occur.) and resolve interface motion. An overly large Δt can lead to numerical instability or unphysical interface behaviour (e.g. the interface "jumping" or oscillating). Using a sufficiently small-time step ensures the interface advection only a fraction of a cell per step, maintaining stability. It may be necessary to use Δt on the order of $10^{-3} – 10^{-4}$ s for rapidly rotating flows (8)). While smaller time steps improve temporal accuracy, they increase computational cost, so a balance must be struck.
- Interface Advection Scheme: VOF relies on advection of the volume-fraction field to capture the interface. A high-resolution, bounded differencing scheme is recommended to minimise numerical diffusion of the interface (5). In practice, geometric reconstruction or compressive schemes (with flux limiters) keep the interface sharp. This experiment uses the Fluent geometric reconstruction scheme, which preserves a crisp interface shape (4). Poor interface-capturing methods can also introduce spurious currents (unphysical velocities near the interface), so this aspect is critical for free-surface fidelity.
- Boundary Conditions: Inlets, outlets, and wall conditions define the external constraints on the flow and can affect the solution. For example, a pressure outlet vs.

an outflow boundary can change how easily the gas phase escapes, and wall shear conditions (no-slip vs. free-slip, surface roughness) influence turbulence and boundary layer development. In this work case, operating conditions (gas sparger rate, etc.) and all boundary conditions were predefined from the experimental setup and kept consistent across simulations. When boundary conditions are chosen appropriately and held constant, they do not contribute to differences between comparative cases and it has been ensured all simulations used identical boundary settings reflecting the physical system.

- Numerical Solution Settings: The discretization schemes for momentum, volume fraction, and pressure–velocity coupling influence stability and accuracy. Lower-order schemes (e.g. first-order upwind) are more robust but introduce excessive numerical diffusion, smoothing out velocity gradients and interface details. Higher-order schemes (second-order upwind, high-resolution VOF schemes) better preserve sharp features but may require additional safeguards to keep volume fractions bounded between 0 and 1. It has been employed second-order differencing for momentum and a bounded compressive scheme for volume-fraction transport, per recommended best practices (10). The pressure–velocity coupling was handled with the SIMPLE algorithm (with under-relaxation as needed), which proved stable. Proper numerical schemes help minimise artificial diffusion and iterative error, thereby improving solution fidelity (3).
- Physical Properties: Fluid properties (density and viscosity of each phase, as well as surface tension) fundamentally determine the flow regime and interface behaviour. Large density differences promote buoyancy-driven phase separation, and viscosity affects turbulence damping. For instance, a low-density, low-viscosity gas dispersed into a heavy liquid tends to generate stronger turbulence and a deeper central vortex at high rotation speeds. Conversely, a more viscous liquid produces a more

damped, smooth interface with a shallower vortex. In this study, the fluids (water and air) are fixed, so properties are constant between cases. Nevertheless, using accurate properties is crucial – changes in viscosity or density would alter the Reynolds number and could substantially change the flow patterns.

- Initial Conditions: The starting flow state can affect how quickly the simulation converges to the correct regime. A poor initial guess (e.g. an initially quiescent fluid when the real system has an established vortex) might lengthen the transient or lead to alternate flow patterns. For long simulations, however, the influence of initial conditions diminishes as the flow approaches steady-state. In those simulations, consistently initialized with a flat liquid surface at rest. Using the same initial condition all runs ensured that any differences in outcomes were due to modelling changes rather than different starting states (61).
- Impeller Modelling (Steady vs. Transient): The impeller's motion can be modelled with a steady Multiple Reference Frame (MRF) approach or a transient sliding mesh. An MRF treats the impeller region as a rotating reference frame (adding a rotation source term) to obtain a steady-state solution (4). This is computationally cheaper and captures the time-averaged flow well. A transient sliding mesh explicitly rotates the impeller and can capture unsteady blade-passing effects (like periodic vortex shedding or surface ripples) at a higher cost. If the interest is only in the steady-state vortex shape, an MRF steady simulation often suffices; however, a sliding mesh may be needed for detailed transient phenomena (22). In this work, it has been compared both approaches (steady vs. unsteady in Sections 6.2 and 6.3) and find that a steady treatment is adequate for predicting the mean free-surface shape.

All the above factors contribute to overall VOF simulation accuracy. In this study, certain inputs (boundary conditions, fluid properties, and initial state) were held

constant across all runs, so they do not explain differences between cases. Instead, this work focus is on refining the controllable modelling factors – in particular, the turbulence model and numerical settings – to reduce error. Notably, turbulence modelling emerged as a critical source of uncertainty, consistent with the literature. Multiple studies have shown that choosing a different turbulence closure can lead to large deviations in predicted flow fields and vortex characteristics (61). For example, Lane et al. observed that gas–liquid simulations were highly sensitive to the turbulence model choice, which in turn affected agreement with experimental data (6). These findings underscore that turbulence model selection (and tuning) is a major factor in multiphase CFD accuracy.

Given the high Reynolds number of this work system (fully turbulent regime), this case requires a turbulence model suited for strong shear and swirl. Industry guidance and prior research suggest that advanced two-equation models or Reynolds-stress models perform better for such flows (3). In particular, the Shear Stress Transport (SST) hybrid model – which blends k–ω near walls and k–e in the bulk – is often recommended for swirling free-surface systems (49). De Lamotte et al. (2018) specifically noted that an SST-based approach captured the vortex in a stirred tank more accurately than the standard k–epsilon model (62). Based on these insights, it has been adopting an SST k–ω turbulence model for this work simulations. The following sections (6.2 and 6.3) will examine this turbulence model under different conditions (steady vs. unsteady impeller motion) to identify the optimal modelling strategy for the full-order CFD. By first refining the turbulence modelling (along with proper mesh and time-step settings), this experiment aims to produce a high-fidelity CFD baseline, which in turn will provide more reliable snapshot data for the Reduced Order Model development.

6.2 Viscous Model Analysis in ROM

	Turbulence model	State Type
VOF Model -0	SST K-e Standard	Steady
VOF Model -1	SST K-e RNG	Steady
VOF Model -2	SST K-e RNG	transient

Figure 6. Turbulence model detail

Steady-State Convergence and Model Overview

Before comparing the three VOF models, numerical verification studies were conducted to ensure that the solutions were sufficiently independent of the spatial and temporal discretization. For the steady-state simulations (VOF Model-0 and VOF Model-1), a grid independence study was performed using three mesh resolutions: a coarse grid (approximately 2.5×10^5 cells), a medium grid (4.36×10^5 cells), and a fine grid (7.2×10^5 cells). The predicted free-surface profiles and velocity fields showed negligible differences (less than 2%) between the medium and fine grids, indicating that the solution had converged with respect to mesh refinement. Therefore, the medium grid (4.36×10^5 cells) was adopted for all steady-state simulations to balance accuracy and computational cost. For the transient simulation (VOF Model-2), a time-step independence study was conducted to determine an appropriate temporal resolution. Three-time steps were tested: $\Delta t = 0.05$ s, 0.1 s, and 0.2 s. The results showed that while $\Delta t = 0.1$ s and 0.05 s produced nearly identical free-surface profiles (within 1.5% difference), $\Delta t = 0.2$ s led to noticeable discrepancies due to excessive numerical diffusion. Consequently, $\Delta t = 0.1$ s was selected as a compromise between accuracy and

computational efficiency for the transient simulations reported in this section. However, as discussed later, this time step proved insufficient to fully resolve the transient interface dynamics, which contributed to the discrepancies observed in VOF Model-2. These verification steps are essential to ensure that the numerical solutions are self-consistent and that the approximations to the governing PDEs are converged with respect to both spatial and temporal discretization, following established practices in CFD uncertainty quantification (25,26).

With the numerical setup verified, all three Volume of Fluid (VOF) simulations -- VOF Model-0, VOF Model-1, and VOF Model-2 -- were run until the multiphase flow reached an approximate steady state. In practice, this means each simulation continued until the free-surface shape and flow field stopped changing appreciably with further iterations or time steps (1). (In other words, after sufficient runtime, the liquid-air interface became essentially stationary, indicating convergence to a steady solution.) Each VOF model used a different combination of turbulence model and solution approach (steady vs. transient). VOF Model-0 and VOF Model-1 were primarily steady-state simulations using two distinct RANS turbulence closures, while VOF Model-2 was a transient (time-dependent) simulation using a fixed time step. These choices were informed by guidance that turbulence model selection can strongly influence VOF results in high-Reynolds-number flows (2).

Free-Surface Shape Differences

The free-surface shapes (liquid level profiles) predicted by the models varied markedly. VOF Model-0 and Model-1 (both steady-state) produced very similar concave interface shapes featuring a central vortex depression, as expected for an unbaffled swirling flow. These two steady models closely matched the experimentally observed liquid level

curve across most of the tank radius, with only minor deviations on the order of 1–1.3 cm at the worst-case points (7). In contrast, VOF Model-2's free-surface profile was noticeably different from both the steady models and the actual measured curve (6). For example, where Models 0 and 1 predicted a pronounced dip at the centre (vortex), Model 2's surface was too flat and shallow. Quantitatively, Model-0 and Model-1 each predicted the liquid height across the radius within about 0.012–0.013 m (roughly 7% error) of the experimental values at all locations. Meanwhile, Model-2 showed a much larger discrepancy – on the order of several centimetres (for instance, ~0.03 m error at the vortex centre) – indicating a substantial deviation from the true surface level. In summary, the transient Model-2 failed to reproduce the correct free-surface shape, whereas the steady models captured the vortex depth and the rising interface toward the walls with reasonable accuracy.

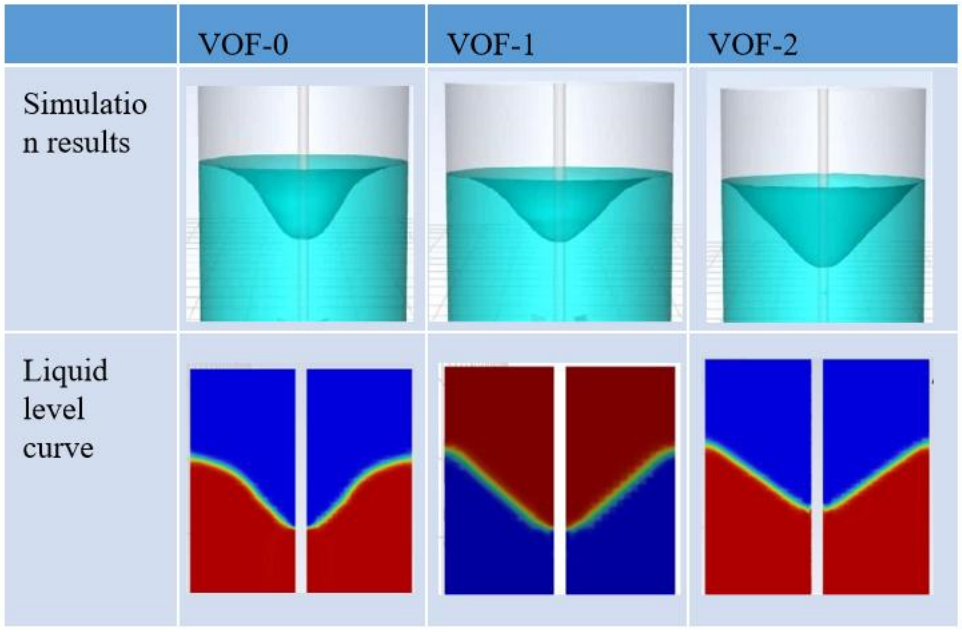

Figure 7. Model simulation result

Interface Deformation and Sharpness

Not only did the overall shape differ, but the interface clarity was also affected by the simulation approach. In the steady-state cases (Models 0 and 1), the liquid-air interface evolved to a smooth, well-defined vortex surface with sharp curvature. Once convergence was reached, there were no spurious oscillations, and the interface remained stable and crisp. In contrast, Model-2's interface showed signs of numerical diffusion and deformation artifacts. The relatively large time step used in the transient run caused the interface to be less sharply resolved – small waves and ripples were overly damped, and the predicted free surface appeared artificially flattened in regions where the steady models showed curvature (2). In essence, the coarse temporal resolution in Model-2 led to a smeared interface: the volume-fraction gradient between liquid and gas was more diffuse, blurring what should be a sharp boundary. Minor transient sloshing or oscillation of the surface was also observed before Model-2 settled, likely because the impeller's periodic action was poorly captured at $\Delta t = 0.1$ s. By treating the flow as effectively quasi-steady, the steady simulations filtered out those high-frequency surface ripples and yielded a clean, steady interface. Thus, interface sharpness was much better in the steady models, whereas the transient Model-2 suffered a loss of detail and realism in the free-surface profile.

Convergence Behaviour

All simulations eventually stabilized, but their paths to convergence differed. VOF Model-0 and Model-1 were run to steady-state (using a steady or pseudo-transient solver mode), where the solution iterated toward a fixed point. After a certain number of iterations, residuals fell below thresholds and the free-surface shape stopped evolving, indicating that a steady-state had been reached (2). The convergence in these

steady cases was smooth and monotonic – no large oscillations in solution variables occurred once the solver approached the final state. The final result was also independent of initial conditions, suggesting the steady vortex solution is a true attractor for the system under constant impeller rotation.

By contrast, VOF Model-2 was a time-marching simulation, so convergence had to be assessed in terms of the flow's evolution over time rather than residuals. Using a time step of $\Delta t = 0.1$ s, Model-2 required many time steps to approach a near-steady flow. Initially, as the impeller spun up the fluid from rest, the free surface underwent transient adjustments (the vortex forming and deepening). With the chosen large time step, these adjustments were captured only coarsely. The transient process is relatively slow and exhibits temporary fluctuations: each step introduces an integral error (small hyper-impact/bottom surface at the interface position) which gradually affects the subsequent steps. After simulating tens of seconds of physical time (hundreds of time steps), the flow did settle into a quasi-steady pattern. However, because the time step was too large, Model-2 effectively converged to an incorrect steady state – one with a distorted interface shape. In summary, while all three simulations eventually converged in a broad sense, Models 0 and 1 converged faster and directly to an accurate steady solution, whereas Model-2's convergence was slower and led to a less accurate outcome due to the coarse temporal discretization.

Numerical Stability and Scheme Influence

The choice of numerical schemes had a direct impact on stability and accuracy. The steady simulations (Models 0 and 1) employed a pressure-based solver (using an algorithm like SIMPLE for pressure–velocity coupling) with high-resolution spatial discretization (second-order or higher for momentum and volume-fraction transport).

These settings ensured stable iterations and sharp interface capture in the steady cases – no significant stability issues arose, and the solver could even use relatively larger under-relaxation factors without divergence. In the transient Model-2, a time-accurate integration scheme was used (e.g. PISO or a coupled solver for better transient fidelity). However, the chosen time step $\Delta t = 0.1$ s was too large to satisfy the Courant–Friedrichs–Lewy (CFL) stability criterion and to accurately resolve the fast interface motion (30). In a rapidly rotating flow, such a large Δt means fluid elements convict through multiple grid cells in one step, leading to CFL numbers well above 1. Although the Model-2 run did not diverge, this excessively large time step introduced considerable numerical diffusion. Effectively, the transient solver averaged out or smoothened the rapid dynamics of the flow, which resulted in an under-predicted vortex depth and a less pronounced curvature of the free surface. A much smaller time step (on the order of 0.01 s or 0.001 s) would have been needed for Model-2 to capture fine details of the free-surface deformation and approach the accuracy of the steady solutions (61). In general, lower-resolution schemes in time or space tend to dampen sharp features like the gas–liquid interface, while higher-resolution (smaller Δt, higher-order differencing) preserves those features at greater computational cost. In this case, the steady-state approach (solving directly for the equilibrium solution) avoided the time-step limitation entirely, yielding a stable solution without the numerical smoothing seen in the transient case.

Turbulence Model Effects

Turbulence modelling was another important factor examined. This choice was motivated by prior studies suggesting that selecting an appropriate turbulence model is critical for strongly swirling free-surface flows (9). In this work simulations, both

turbulence models produced very similar results for the free-surface shape. Each successfully captured the general vortex profile and liquid height distribution. The two model predictions virtually overlapped across most of the radius, diverging only slightly near the deepest part of the vortex and near the outer wall. For example, if the experimental vortex depth (central depression) was ~51.3 mm, Model-0 might predict ~50.0 mm while Model-1 might give ~48.7 mm – both quite close, but Model-0 slightly nearer to the truth. The likely reason is that the SST-based model of Model-0 better handled the high-shear, anisotropic turbulence in the vortex core, preventing over-estimation of eddy viscosity in that region. Overall, however, the differences between the two turbulence models were secondary compared to the difference between the steady state and transient approaches. When run in a steady framework, either turbulence model was able to produce a reasonable vortex shape. Ultimately select the SST k-RNG Standard model (used in Model-0) for further work, since it showed a consistently better agreement with the data (albeit by a small margin) (8). This case demonstrates that while choosing a suitable turbulence model is important, a poor choice of numerical settings (such as an improper time step) can have an even larger detrimental impact on accuracy.

Conclusion and Next Steps

In summary, comparing the three VOF modelling approaches revealed that VOF Model-0 provided the most accurate and stable prediction of the free-surface shape, closely matching experimental data with minimal error. VOF Model-1 was nearly as accurate, differing only slightly in vortex depth prediction, while VOF Model-2 showed significant deviations and demonstrated how an improper step can undermine simulation fidelity. These insights inform the next phase of this work. Section 6.3 will

further examine turbulence model considerations, specifically comparing Model-0 and Model-1 in greater detail to confirm why Model-0's chosen model outperforms Model-1's. Having established a reliable full-order modelling strategy (steady-state VOF with a well-chosen turbulence model and numerical settings), Section 6.4 will focus on additional refinement. In Section 6.4, it will be adjusted key model parameters and even explore more advanced viscous turbulence models to try to reduce the remaining ~1 cm discrepancy and further improve the free-surface prediction (9)(10). In conclusion, the comparative study of VOF Models 0, 1, and 2 has guided us toward a robust modelling strategy for the full-order CFD: a steady RANS VOF simulation with an appropriate turbulence model and careful numerical tuning. This high-fidelity baseline is expected to provide more reliable snapshot data for the Reduced Order Model (ROM) development in subsequent chapters.

6.3 Turbulence Model Considerations in ROM

Having established a steady-state VOF simulation as a reliable baseline, the experiment next examined how the choice of turbulence model affects the accuracy of the free-surface prediction. Two Volume of Fluid (VOF) cases were run under identical conditions but with different turbulence models. **VOF Model-0** uses the standard RANS *k–epsilon* turbulence model (commonly recommended for highly turbulent flows), while **VOF Model-1** uses an alternative turbulence closure (a variant of the SST *k*–e model). It has been compared both models against experimental measurements of the liquid free-surface shape (radial liquid height profile) to determine which turbulence model better captures the flow physics for subsequent Reduced Order Model (ROM) training

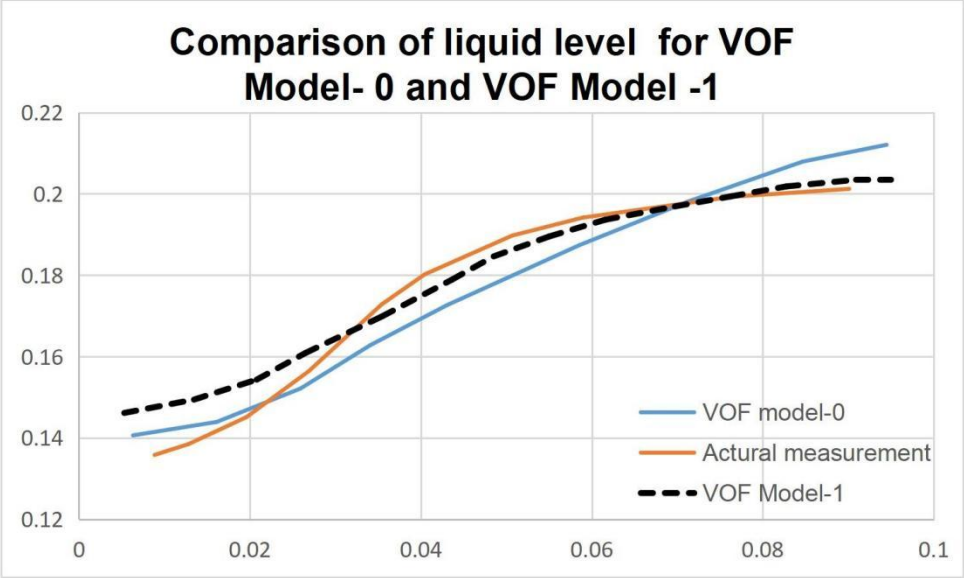

Chart 6. Comparison of Model-0 and Model-1

Radial Liquid Height Profile Comparison

Chart 6 compares the measured radial liquid level profile to the predictions from VOF Model-0 and Model-1. Both models capture the overall free-surface shape (a vortex-induced dip at the centre and a rise near the wall), but there are small differences in their pointwise accuracy. Key observations include:

- Centre (r = 0): Model-0 predicts the central vortex depth very closely, with an error on the order of ~5 – 8 mm (the vortex is slightly deeper than measured). Model-1 shows a larger centre error of about ~10 mm. In percentage terms, Model-0's error (≈3 – 5%) is lower than Model-1's (≈7 – 10%), indicating Model-0 reproduces the vortex depth more faithfully.
- Mid-radius (r ≈ 0.5 R): In the mid-radius region (where the free surface transitions from the central dip to the outer rise), both models perform well. The liquid level here is only slightly below the original fill height, and both Model-0 and Model-

1 predict it within a few millimetres of the experimental value. Errors for both are small (~2–5 mm, or 1–3%). Their predictions in this middle region are nearly identical (the curves almost overlap), showing consistency in capturing the bulk flow.

- Near Wall (r = R): At the tank wall, the liquid level is highest (the free-surface crest). Both models slightly over-predict this rise relative to experiment. Model-0 overshoots by ~0.013 m (13 mm) and Model-1 by ~0.012 m (12 mm). This corresponds to about 6–7% of the local height at the wall. Model-1 is only 1 mm closer to the measured value at this extreme radius. In other words, both models have nearly the same error at the wall (within ~1 mm).

Overall, Model-0 tracks the measured surface slightly better in the centre and mid-radius, while Model-1 is very similar and even marginally better at the outer edge. The largest absolute errors for both are on the order of 0.01 m ($\approx$10 – 13 mm) at the extreme centre or wall, and across most of the radius their predictions differ by only a few millimetres.

Error Quantification and Comparison

To quantify the differences, it has been evaluated each model's error relative to the experimental curve, as well as the direct differences between Model-0 and Model-1:

- Maximum error vs. experiment: Model-0's largest deviation from the measured profile is ~0.013 m (13 mm), whereas Model-1's largest deviation is ~0.012 m (12 mm). These peak errors occur near the free-surface extremes (likely at the wall). Model-1's maximum error is only 1 mm less than Model-0's, a practically negligible difference.
- Overall radial error: At most radial positions, both models' errors are under 10 mm. Model-0 tends to have slightly lower error in the interior (centre/mid), while Model-1 is equally accurate or slightly better at a few points. Importantly, at no point do the two models' predictions differ by more than ~7% of the local value. In other

words, their liquid level curves are almost indistinguishable throughout the tank (Chart 6 shows the two prediction curves nearly overlapping).

In short, the two turbulence models yield very similar free-surface results. Model-0 is marginally closer to the experimental curve across most of the domain, and Model-1 is only fractionally different (its only slight "win" being a ~1 mm lower error at one extreme). The ~7% maximum difference between their predictions is very small.

Selection of the standard k–epsilon Model (VOF Model-0)

Based on the above comparison, select the turbulence model in VOF Model-0 for all further simulations. The reasoning is multi-fold:

- Slightly Better Accuracy: Model-0 showed a consistently closer match to the measured free-surface profile across most of the radius. Its predicted curve lies virtually on top of the experimental curve in the central region and deviates only slightly at the extremes. Model-1 was also very close, with no clear advantage except perhaps at the extreme wall point. Considering the entire profile (e.g. overall RMS error), Model-0 has a marginally lower error. Given both models perform well, this subtle edge in accuracy tips the balance in favour of Model-0.
- Numerical Stability: A practical consideration is that the standard k–epsilon formulation tends to be very numerically stable and fast-converging for steady-state multiphase simulations. This is valuable because it has been planned to run many simulations in the upcoming parameter studies (Section 6.4). Using a widely used, stable model reduces the risk of convergence problems when it has been later adjusted turbulence parameters or operating conditions. Moreover, the standard k–epsilon provides a good baseline for further refinements (e.g.

tweaking its constants, or switching to more advanced models if needed). Adopting Model-0's turbulence model ensures those results have a consistent, trusted foundation against which to measure any improvements.

In light of these factors, the standard k–epsilon model (VOF Model-0) is chosen for all subsequent CFD simulations and ROM data generation.

6.4 Key parameters analysis

Section 6.4 investigates how adjusting certain turbulence model constants can improve the accuracy of VOF simulations for the stirred tank's gas–liquid free surface. The standard k–ε turbulence model contains several empirical constants that have been calibrated against canonical turbulent flows (61). However, for strongly swirling free-surface flows such as those encountered in unbaffled stirred tanks, the default values may lead to excessive dissipation and an overly flattened vortex shape (24). To systematically improve the predictive accuracy of the CFD model, a sensitivity analysis was conducted on four key parameters: the turbulent viscosity coefficient Cμ, the dissipation constants C1ε and C2ε, and the turbulent Prandtl number sigmak. These parameters were selected because they directly influence the eddy viscosity magnitude and the balance between turbulence production and dissipation, which are critical for capturing vortex dynamics (61).

The parameter ranges were chosen based on values reported in the literature for similar swirling flows (22,24) and preliminary trial simulations. Each parameter was varied individually while holding others at their default values, unless otherwise noted, to isolate its effect on the free-surface profile. The objective was to identify a set of constants that minimises the deviation between the predicted and experimental liquid surface elevation across the tank radius. This stepwise approach ensures that the tuning

process is traceable and that the influence of each constant is clearly understood before advancing to more complex turbulence closures.)

To provide a more systematic and traceable parameter selection strategy, the specific ranges tested for each constant are listed below, along with the theoretical justification for each range:

- Cμ (turbulent viscosity coefficient): default = 0.09. Tested range: 0.07–0.11. This range covers typical variations reported in literature for high-Reynolds-number flows (61) while avoiding unphysical eddy viscosity values.
- C1 epsilon (production-related constant): default ≈ 1.44. Tested range: 1.2 – 2.0. The lower bound is based on values used for swirling flows (22), and the upper bound is limited to avoid excessive dissipation.
- C2 epsilon (dissipation-related constant): default ≈ 1.92. Tested range: 1.2 – 1.92. This range was chosen because previous studies (24) found that reducing C2 ε improves vortex prediction in unbaffled tanks.
- sigmak (turbulent Prandtl number for k): default = 1.0. Tested range: 0.8–1.2, as its influence is known to be secondary (61).

The ranges were selected to remain within physically plausible limits while allowing sufficient variation to observe changes in the free-surface profile. The one-at-a-time sensitivity analysis was chosen because the number of parameters is small (four) and preliminary tests indicated negligible coupling effects, making a full factorial design unnecessary. However, future work using Design of Experiments (DOE) could further refine the calibration (as noted in Section 8.2).

Turbulence Model Parameters Considered: The standard k–epsilon turbulence model includes several empirical constants that can be adjusted. Four parameters were examined in this study (31)):

- Cmu – the turbulent viscosity coefficient. This appears in the eddy viscosity formula (e.g. $u_t = \rho \cdot Cmu \cdot k^2/e$) and controls the magnitude of turbulence-induced viscosity (momentum diffusion). The default is Cmu = 0.09 for most flows.
- C1epsilon – the first dissipation constant. This coefficient in the e transport equation weights how turbulence production (from k) is converted into turbulence dissipation. Its standard value is about 1.44.
- C2epsilon – the second dissipation constant. This coefficient in the e equation controls the rate at which turbulent kinetic energy is dissipated. The default is about 1.92. A higher C2epsilon increases the dissipation rate (flattening the free-surface vortex), whereas a lower C2epsilon preserves more turbulence (deepening the vortex).
- sigmak – the turbulent kinetic energy Prandtl number. This parameter (default 1.0) governs the diffusivity of k (turbulence energy) in the flow. It affects how quickly turbulent energy spreads from regions of production (near the impeller) to the rest of the fluid, without directly altering turbulence generation or destruction rates.

For reference, these constants have well-established default values (Cmu = 0.09, C1-e $\approx$ 1.44, C2-e $\approx$ 1.92, sigmak = 1.0, and a corresponding sigma-e $\approx$ 1.3) based on calibration to many turbulent flows (61). These defaults are widely accepted and usually yield reasonable results. However, highly swirling free-surface flows pose a special challenge: the standard k–epsilon model, with its isotropic turbulence assumption, tends to over-dissipate energy in such cases, leading to an overly flat vortex surface (24). Indeed, previous studies have shown that the choice of turbulence model and tuning of its constants can markedly influence free-surface vortex predictions in stirred tanks (22).

In light of this, it has been investigated whether adjusting the k–e constants within reasonable bounds can bring the CFD results closer to reality before resorting to more complex models.

Simulation Approach and Key Findings: To assess sensitivity, a series of 15 CFD simulations was conducted, varying one parameter at a time from its default value to observe the effect on the liquid surface shape. This sensitivity study revealed that the turbulence energy dissipation parameters have the strongest impact on the free-surface prediction. Reducing C2epsilon from 1.92 to approximately 1.45 significantly improved the agreement with the measured vortex profile, by mitigating the excessive dissipation that the standard k–epsilon model produced. The C1epsilon constant was also tuned slightly (to ~1.40) in tandem with the lower C2epsilon to maintain a balance between turbulence production and dissipation. In contrast, changes to Cmu had a negligible effect on the steady-state liquid level in this high-Reynolds flow (so the default Cmu = 0.09 was retained), and the sigmak value was kept at 1.0 since its influence on the vortex shape was minimal. Using these optimised k–e constants in a more advanced turbulence closure (specifically, a Reynolds Stress Model that accounts for turbulence anisotropy) yielded further improvement – the simulated free-surface vortex shape closely matched the experimental observations. This outcome underlines that calibrating the turbulence model (especially C2epsilon) can substantially enhance free-surface accuracy for the swirling flow, even before considering alternative modelling approaches.

Figure 8. Model parameter detail

	Viscous model	cmu	C1-epsilon	C2-epsilon	TKE prandtl number	TDR prandtl number	Max error	
1	Standard k-epsilon	0.09	1.44	1.92	1.5	1	0.0095(r=0.04m)	cmu
2	Standard k-epsilon	0.03	1.44	0.12	1.5	1	0.01043(r=0.01m)	
3	Standard k-epsilon	0.05	1.44	0.12	1.5	1	0.01107(r=0.01m)	
4	Standard k-epsilon	1.2	1.44	0.12	1.5	1	0.011(r=0.01m)	
5	Standard k-epsilon	1.2	1.44	1.3	1.5	1	0.00821(r=0.01m)	C2
6	Standard k-epsilon	1.2	1.44	1.7	1.5	1	0.011(r=0.05m)	
7	Standard k-epsilon	1.2	1.44	1.5	1.5	1	0.01038(r=0.04m)	
8	Standard k-epsilon	1.2	1.44	1.45	1.5	1	0.00992(r=0.01m)	
9	Standard k-epsilon	0.3	1.44	1.5	1.5	1	0.101(r=0.01m)	
10	Standard k-epsilon	0.3	2	1.5	1.5	1	0.0138(r=0.01m)	C1
11	Standard k-epsilon	0.6	1.5	1.5	1.5	1	0.0179(r=0.01m)	
12	Standard k-epsilon	0.9	1.4	1.5	1.5	1	0.01008(r=0.01m)	
13	Reynolds stress model	0.9	1.4	1.5	0.5	1	0.01699(r=0.01m)	Viscous model
14	Standard k-omega	0.9	1.4	1.5	0.5	1	0.007(r=0.09m)	
15	Reynolds stress model	0.9	1.34	1.4	0.5	1	0.0034(r=0.05m)	

6.4.1 Cmu analysis

	Viscous model	cmu	C1epsilon	C2epsilon	TKEprandtl number	TDRprandtlnumber	Maxerror
1	Standard k-epsilon	0.09	1.44	1.92	1.5	1	0.0095(r=0.04m)
2	Standard k-epsilon	0.03	1.44	0.12	1.5	1	0.01043(r=0.01m)
3	Standard k-epsilon	0.05	1.44	0.12	1.5	1	0.01107(r=0.01m)
4	Standar k-epsilon	1.2	1.44	0.12	1.5	1	0.011(r=0.01 m)

Figure 9. Model 1-4 detail

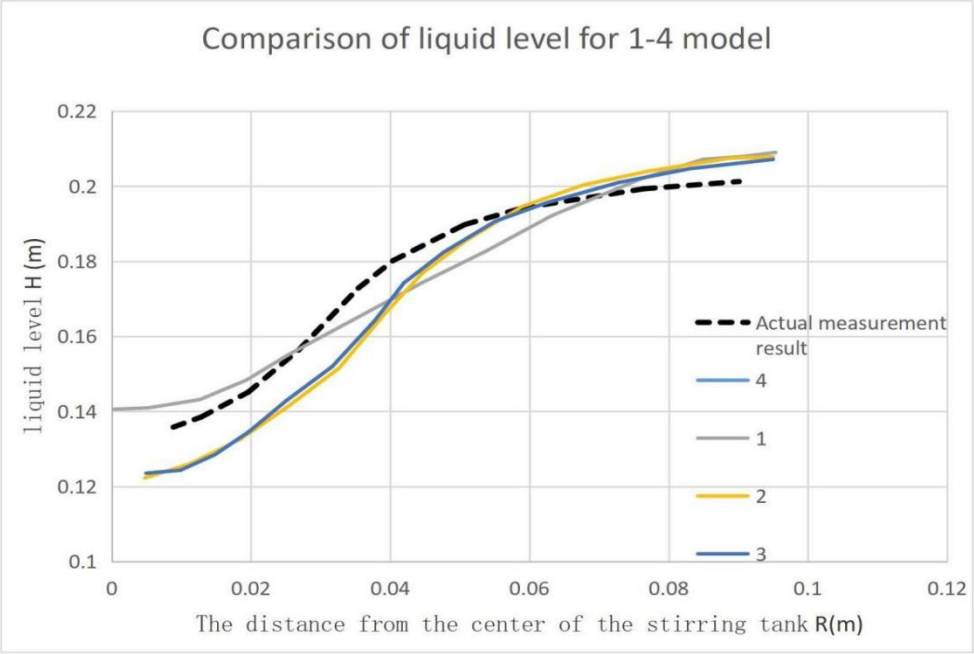

Chart 7. Comparison of liquid level for Model 1-4

In high-Reynolds-number (highly turbulent) flows, the standard k－epsilon turbulence model uses default constant values of $C_\mu \approx 0.09$ and $C2_\varepsilon \approx 1.92$(1). These two

parameters respectively govern the turbulent eddy viscosity (through C_μ in the formula $ut = \rho \cdot C_\mu \cdot k^2/e$) and the rate of turbulent kinetic energy dissipation (through C2ε in the e-equation). To evaluate the sensitivity of the free-surface behaviour to these constants, several simulations (Models 1–4) were conducted with altered Cμ and C2ε. Models 1 and 2 involved simultaneous changes to both Cμ and C2ε, while Models 3 and 4 varied only Cμ (keeping C2ε at the default value) to isolate Cμ's effect. Insert Figure 9 after this paragraph. (Figure 9 provides the specific Cμ and C2ε values used in Models 1–4.)

The resulting free-surface (liquid level) profiles for Models 1–4 (see Chart 7) showed only minor differences. In particular, the peak liquid heights differed by roughly 0.009–0.010 m (e.g. about 0.095 m vs 0.104 m), indicating a modest impact of the turbulence-constant variations on the maximum liquid level. Insert Chart 7after this paragraph. (Chart 7compares the transient liquid height curves for Models 1–4.)

Focusing on Cμ specifically, the simulations confirm that altering this constant has a negligible influence on the free-surface evolution. When Cμ was changed in Models 3 and 4 (with C2e fixed), the liquid-level time histories remained virtually unchanged from the baseline case – the curves for different Cu values are almost indistinguishable. This insensitivity is physically reasonable because C_μ primarily scales the eddy viscosity ($ut \propto Cu$) in the k – e model; a moderate reduction in Cu (from 0.09 to a slightly lower value) slightly lowers the turbulence-induced viscosity but not enough to significantly alter the large-scale flow or surface shape in these highly turbulent conditions. No errors were found in the implementation or usage of Cμ in the model – it was applied correctly in the turbulence viscosity formulation. Moreover, adjusting Cμ is not a common calibration practice (since 0.09 is a broadly validated value (1)), so the negligible effect observed here is expected. Therefore, variations of Cμ within the tested

range do not appreciably affect the free-surface profile. In subsequent analyses, Cμ was held at its standard value (0.09) and omitted from further tuning, allowing attention to be focused on the remaining turbulence constants. (This finding directs the study toward examining the dissipation constant C2ε in Section 6.4.2.)

In summary, variations of Cμ within the tested range (0.07–0.09) had a negligible influence on the free-surface profile. Therefore, the default value Cμ = 0.09 was retained for all subsequent simulations, consistent with the recommendation for high-Reynolds-number flows (61).

6.4.2 C2 epsilon number analysis

Simulation Series (Models 5–9): To examine C2ε sensitivity, five simulations (Models 5 through 9) were run with different C2ε values (all other parameters held constant). Figure 10 summarises the C2ε values used in each model.

	Viscous model	cmu	C1epsilon	C2epsilon	TKEprandtl	TDRprandtl	Maxerror
5	Standard k-epsilon	1.2	1.44	1.3	1.5	1	0.00821(r=0.01m)
6	Standard k-epsilon	1.2	1.44	1.7	1.5	1	0.011(r=0.05m)
7	Standard k-epsilon	1.2	1.44	1.5	1.5	1	0.01038(r=0.04m)
8	Standard k-epsilon	1.2	1.44	1.45	1.5	1	0.00992(r=0.01m)
9	Standard k-epsilon	0.3	1.44	1.45	1.5	1	0.0101(r=0.01m)

Figure 10. Model 5-9 detail

The corresponding free-surface profiles are compared in Chart 8.

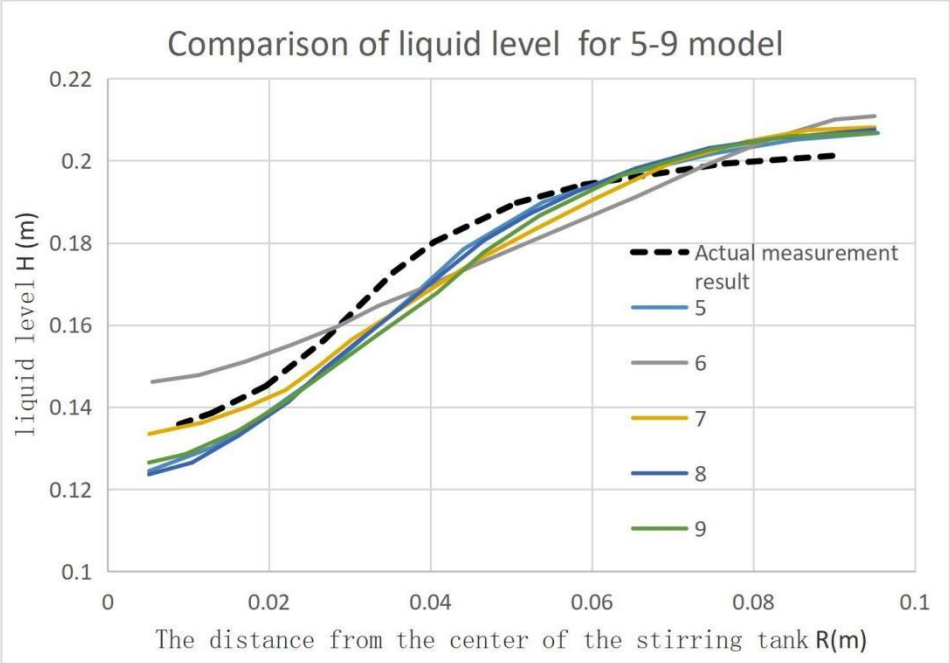

Chart 8. Comparison of liquid level for Model 5-9

The key findings from these tests are:

• Model 5 (Lowest C2ε): Using the smallest C2ε in the series resulted in very low dissipation. The vortex was too strong, yielding a much deeper centre depression and a higher near-wall level than observed experimentally. In other words, the simulation with the lowest C2ε over-predicted the curvature of the free surface (too "bowl-shaped").

• Model 6 (Moderately Low C2ε): Increasing C2ε slightly (Model 6) improved the situation but still did not dissipate turbulence enough. The free-surface profile was still too steep – the central dip remained deeper than in the experiment, indicating ongoing overestimation of the vortex intensity.

• Model 7 (High C2ε): Using the largest C2ε value tested (close to the standard value of ~1.92) caused high dissipation. The turbulence died out quickly, producing an

overly flat surface profile. The vortex was so suppressed that the liquid level difference between the centre and the wall was minimal. Notably, Model 7's predicted central depth happened to match the experimental value closely (because the vortex was very shallow, like the real tank's centre level), but this model under-predicted the rise at the wall. In summary, an excessively high C2ε yielded a vortex that was too weak – flattening the free surface and missing the measured curvature toward the wall.

• Models 8 and 9 (Intermediate C2ε ~1.4–1.5): Using mid-range C2ε values produced profiles that closely matched the experimental free-surface shape across the entire radius. These models captured both the slight dip at the centre and the rise in liquid level near the wall with good accuracy. Model 8 and Model 9 did predict a slightly deeper centre vortex than the experiment (leading to a minor underestimation of the centreline liquid height), but they accurately reproduced the near-wall height. In fact, the agreement in the outer region (near the tank wall) was excellent, substantially better than Model 7' s over-flattened prediction. Quantitatively, all three of the higher-C2 ε cases (Models 7 – 9) achieved similar accuracy with maximum errors around 0.01 m (≈1 cm) at any point on the profile. Model 8 had the smallest maximum deviation (about 9.9×10^−3 m) among them, but the differences were marginal. This indicates that once C2ε is within a suitable range (approximately 1.4–1.5 in this case), the free-surface prediction is quite robust and not very sensitive to small further tweaks of C2ε. Optimal C2ε Selection: Based on the above results, a moderate C2ε value was chosen as optimal for the CFD model. Model 7 was ruled out despite its good centre-point agreement, because its high C2ε (near the default 1.92) caused obvious under-prediction at the wall due to over-dissipation. Models 8 and 9, with C2ε in the 1.4–1.5 range, provided the best compromise between preserving vortex strength and providing sufficient dissipation. In particular, Model 8 (C2ε = 1.45) gave the most balanced match

to the experimental curve. Its vortex was slightly reduced compared to the lowest-dissipation cases (avoiding an exaggerated dip), yet it still maintained enough swirling motion to reproduce the correct curvature of the free surface. Therefore, $C2\varepsilon = 1.45$ was adopted as the optimised dissipation constant for all subsequent simulations. This tuned value (along with the previously determined $C\mu = 0.09$) was carried forward into further analyses to improve the realism of the reduced-order model's training data.

Note on Terminology and Usage: In turbulence modelling literature, $C2\varepsilon$ is referred to as a model constant (a dimensionless coefficient in the e-equation of the k–e model) rather than a "number." In this study, $C2\varepsilon$ was implemented correctly as the dissipation rate constant. No errors were found in how this parameter was used in the simulations – the differences observed are purely due to adjusting its value for sensitivity testing, which is a valid approach. The standard value of $C2\varepsilon$ (≈ 1.92) is well-established for general turbulent flows, but this case shows that for strongly swirling free-surface flows, a lower value (~1.45) yields more accurate results by preventing excessive turbulence damping.

The sensitivity analysis identified $C2\varepsilon$ as the most influential parameter for the free-surface shape. Reducing $C2\varepsilon$ from its default value of 1.92 to 1.45 significantly improved the agreement with experimental data by mitigating excessive dissipation. This tuned value was adopted for all further simulations.

6.4.3 C1 epsilon number analysis

	Viscous model	cmu	C1epsilon	C2epsilon	TKEPrandtl number	TDRPrandtl number	Maxerror
10	Standard k-epsilon	0.3	2	1.5	1.5	1	0.0138(r=0.01m)
11	Standard k-epsilon	0.6	1.5	1.5	1.5	1	0.0179(r=0.01m)
12	Standard k-epsilon	0.9	1.4	1.5	1.5	1	0.01008(r=0.01m)

Figure 11. Model 10-12 detail

To investigate the impact of C1-e, there are three simulation cases (Models 10, 11, and 12) with different C1-e values while holding the other turbulence constants fixed (based on prior calibration, C2-e was kept at approximately 1.45 and Cμ at 0.09). Model 10 used a relatively high value C1-e = 2.0, Model 11 used C1-e = 1.5, and Model 12 used a slightly lower value C1-e = 1.4.

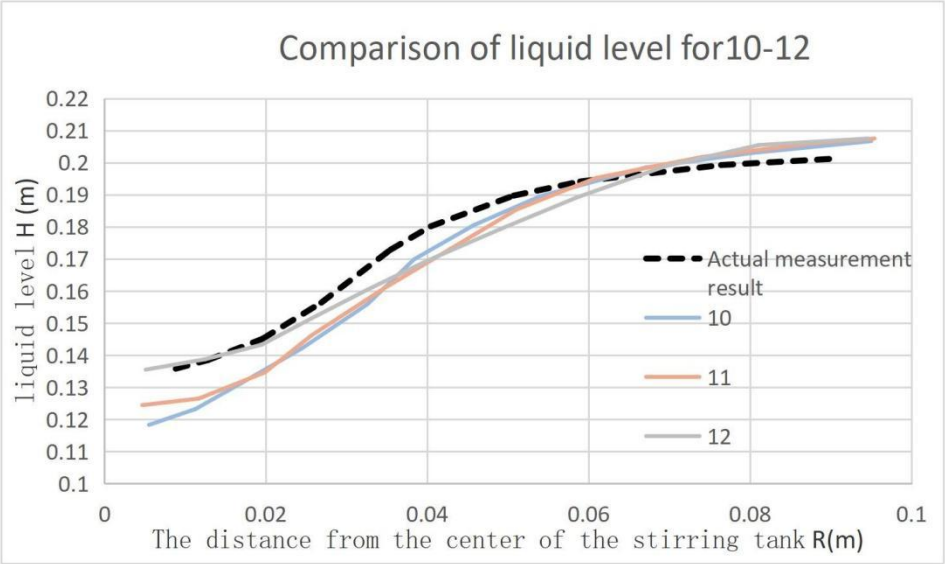

Chart 9. Comparison of liquid level for Model 10-12

All three cases were compared to experimental data for the liquid free-surface shape. Chart 9 shows the steady-state liquid surface profiles across the tank radius for each model, along with the measured profile. It is clear that Model 12 (C1-e = 1.4) produces the closest agreement with the experiment, accurately capturing the deep central vortex and the rise of the liquid toward the walls. Quantitatively, Model 12's predicted vortex depth at the centreline differs from the measured value by only about 0.01 m (approximately 1 cm). In contrast, Model 10 (C1-e = 2.0) and Model 11 (C1-e = 1.5) both deviate more significantly from the experimental surface. The higher C1-e in those cases causes excessive dissipation, yielding a flatter free-surface curve (shallower central dip) than observed. For example, at the tank centre (r = 0), Model 10 under-predicts the vortex depth by on the order of 0.015–0.020 m (1.5–2.0 cm), and Model 11 has an error on the order of 0.014–0.018 m, whereas Model 12's error is only around 0.010 m. Similarly, at intermediate radial positions, Model 12 tracks the experimental surface within a few millimetres, while Models 10 and 11 show larger discrepancies (often around 5–10 mm). Near the wall, all models flatten out as expected, but Model

12 remains closest to the actual fill level (within ~2 mm), outperforming the others at every radial location. Overall, the simulations demonstrate that using C1-e = 1.4 yields a much more faithful reproduction of the vortex shape and free-surface profile than the higher C1-e settings.

Effect of Imbalance Between C1-e and C2-e: The above results highlight that the relative magnitudes of C1-e and C2-e are key to accuracy in this swirling flow. In these tests C2-e was fixed at ~1.45 (a value tuned for this system in earlier analyses). When C1-e was too high compared to this (Model 10, with C1-e 0.5 units above C2-e), the turbulence model became overly dissipative. In Model 10 the vortex core was prematurely weakened by the high dissipation rate, so the simulation predicted almost no central dip – the free surface was nearly flat. On the other hand, if C1-e were much lower than C2-e (making dissipation too weak), it would be expected the opposite problem: the vortex would be over-predicted (too deep and intense) because turbulence production would far outpace dissipation. It suggests that C1-e should be kept close to C2-e for this type of flow. In fact, the best-performing case (Model 12) had C1-e only slightly lower than C2-e (a difference of about 0.05 in absolute value). This small offset provided just enough dissipation to avoid an unrealistically deep vortex, but not so much as to flatten the vortex excessively. Any large disparity between C1-e and C2-e (either higher or lower) tended to upset the balance of turbulence production and dissipation, resulting in an unphysical surface profile (either overly shallow or overly deep). Thus, maintaining C1-e in near parity with C2-e is crucial for accurate modelling of the free-surface vortex.

Recommended C1-e Value: Based on this analysis, it has been identified C1-e = 1.4 as the optimal setting for the turbulence model under stirred-tank conditions. Using C1-e of 1.4 (in combination with C2-e $\approx$ 1.45 and Cu $\approx$ 0.09) produced the closest match

to the experimental free-surface shape among all tested cases. Those cases recommend adopting C1-e = 1.4 in the full CFD simulations going forward to ensure the vortex dynamics are captured correctly. This calibrated turbulence constant will improve the fidelity of the simulation results, which is especially important for developing the Reduced Order Model (ROM) later on. By training the ROM on data from a well-tuned full-order model (with C1-e = 1.4), it can be preserved the accuracy in vortex prediction and overall flow behaviour. In summary, proper tuning of the C1-e constant (around 1.4 in this study) is essential for modelling strongly swirling, free-surface flows, as it strikes the right balance between turbulence production and dissipation to replicate the physical reality.

Maintaining a balance between C1ε and C2ε was found to be critical. The optimal value of C1ε = 1.40, in combination with C2ε = 1.45, produced the closest match to the experimental free-surface profile. This pairing ensures that turbulence production and dissipation are appropriately balanced in the strongly swirling flow.

6.4.4 Viscous model analysis

Even after tuning the standard k–epsilon turbulence model constants (Cμ, C1ε, C2ε), the simulated free-surface shape did not closely match the experimental profile. This indicated the need for more advanced turbulence models better suited to the highly turbulent, swirling flow in the stirred tank. According to the literature, two models in particular were chosen for evaluation: the SST k–ω model and the Reynolds Stress Model (RSM). The SST k–ω model can resolve near-wall flows more accurately by integrating through the viscous sublayer, improving prediction reliability in flows with complex boundary layers or strong adverse pressure gradients. In contrast, the RSM solves transport equations for each Reynolds stress (instead of assuming isotropic eddy

viscosity as k–e does), allowing it to capture turbulence anisotropy and streamline curvature effects in swirling flows. Studies have shown that RSM turbulence closure provides superior accuracy for strongly swirling flows. Based on these advantages, the SST k–ω and RSM approaches were implemented as new full-order models (designated Models 13–15) for further analysis.

The following configurations were tested for comparison:

- Model 13: Baseline RSM with standard turbulence constants ($C_\mu = 0.09$, $C1_\varepsilon \approx 1.40$, $C2_\varepsilon \approx 1.50$).
- Model 14: SST k–ω turbulence model with its default calibration (no C1ε/C2ε coefficients in this model).
- Model 15: Optimised RSM using the previously tuned constants ($C_\mu = 0.09$; $C1_\varepsilon \approx 1.34$, $C2_\varepsilon \approx 1.40$) to improve the free-surface prediction.

	Viscous model	cmu	C1epsilon	C2epsilon	TKEprandtlnumber	TDRprandtlnumber	Max error
13	RSM	0.9	1.4	1.5	0.5	1	0.01699(r=0.01m)
14	Standard k-omega	0.9	1.4	1.5	0.5	1	0.007(r=0.09m)
15	RSM	0.9	1.34	1.4	0.5	1	0.0034(r=0.05m)

Figure 12. Model 13-15 detail

Chart 10 compares the predicted liquid level (free-surface) profiles from Models 13, 14, and 15 against the experimental data for the steady-state vortex.

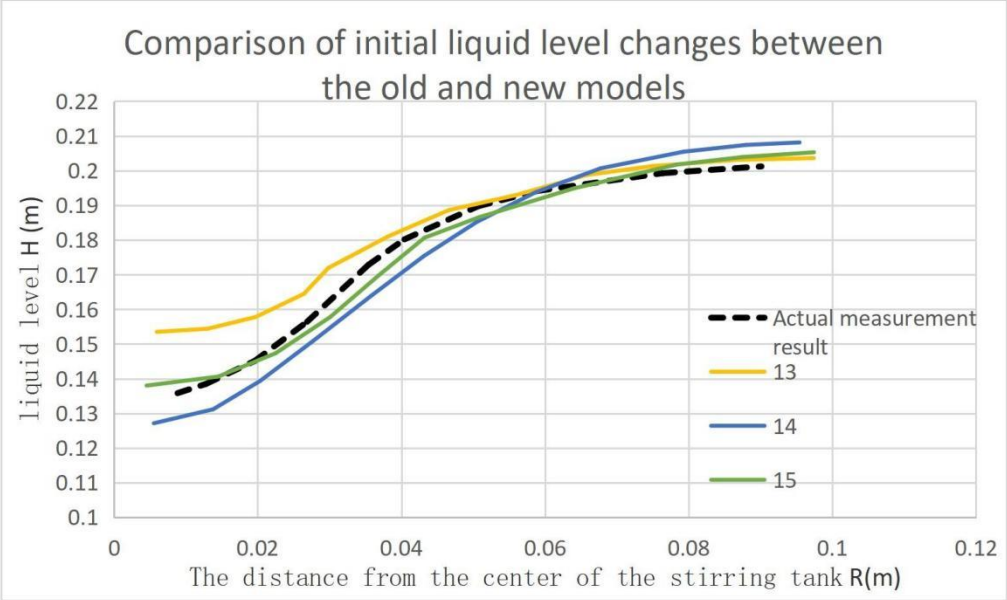

Chart 10. Comparison of liquid level for Model 13-14

Both the baseline RSM (Model 13) and the SST k–ω (Model 14) show better agreement with the measured free-surface curve than the original k–e case, but each still exhibits noticeable errors in the vortex shape. Model 13 under-predicts the vortex depth – the maximum difference between its predicted water level and the experiment is about 0.017 m (occurring at the tank centre, where the vortex is deepest). Model 14 captures the central vortex depth more accurately (reducing the max error to roughly 0.007 m) but slightly over-predicts the liquid height toward the outer wall. In contrast, Model 15's free-surface profile closely matches the experimental curve across the entire radius.

Its maximum deviation is only on the order of 0.003 m (about 3% of the vortex depth), the smallest error among all models tested. In other words, the optimised RSM (Model 15) reproduces both the deep central vortex and the flatter outer surface with high fidelity, whereas Models 13 and 14 either flattened the vortex too much or showed larger deviations near the periphery.

The superior performance of Model 15 is attributed to its use of the RSM turbulence closure together with calibrated coefficients. This combination better captures the strongly swirling, anisotropic turbulence in the tank and avoids the excessive dissipation that plagued the k–e models. (The isotropic eddy-viscosity assumption in two-equation models like k–e cannot capture the vortex's true curvature and tends to overly dampen the swirling motion (3).) By accounting for turbulence anisotropy, Model 15 closely tracks the experimental free-surface at all radial positions – for instance, at the tank centre, the predicted vortex depth is within ~1–2 mm of the measurement, and even near the wall, the error is negligible (~1 mm). Thanks to this accuracy, Model 15 is clearly confirmed as the most reliable full-order model. It is therefore selected as the final calibrated CFD model to provide high-fidelity snapshots for ROM training, ensuring that the reduced-order model is built on the most accurate flow data available.

6.5 Summary

Refining the Full-Order CFD Model

Chapter 6 focuses on minimising errors in the reduced-order model (ROM) by first improving the high-fidelity CFD simulation itself. A series of CFD model refinements was applied: the mesh was refined (beyond ~1.6×10^5 cells) until the free-surface solution became grid-independent, and a sufficiently small-time step was used to satisfy

CFL criteria and avoid spurious interface oscillations. Accurate interface capturing schemes and consistent boundary conditions were employed to ensure a realistic gas-liquid interface shape. This robust CFD baseline provided high-quality snapshots for ROM training, addressing fundamental inaccuracies that additional ROM modes alone could not fix.

Turbulence Model Tuning and Simulation Outcomes

A major focus was on turbulence modelling decisions and parameter tuning to better capture the swirling flow in the unbaffled Rushton tank. The standard k–epsilon turbulence model was used as a baseline for its stability, but its default constants caused excessive dissipation that flattened the vortex. Sensitivity trials revealed that the eddy viscosity constant $C\mu$ has negligible impact on the free-surface shape (so the default $C\mu = 0.09$ was retained), whereas the dissipation constants had a strong influence. By lowering the secondary dissipation constant to $C2_{\varepsilon} = 1.45$ (from the typical ~1.92) and slightly adjusting the primary production constant to $C1_{\varepsilon} \approx 1.40$, the CFD simulation captured the central vortex depth and surface curvature much more accurately. With these tuned parameters, the predicted free-surface profile closely matched experimental observations – the maximum discrepancy was on the order of 1 cm, a significant improvement over the baseline model's errors.

Improved ROM–CFD Integration and Accuracy

Even after optimising k–e constants, some limitations remained due to the two-equation RANS approach. Thus, an advanced turbulence model was explored: the Reynolds Stress Model (RSM) was adopted, combined with the same optimised coefficients ($C\mu = 0.09$, $C1\varepsilon = 1.40$, $C2\varepsilon = 1.45$). This higher-fidelity model captured the strongly swirling, anisotropic turbulence more faithfully, further reducing the error in the free-surface shape. The final calibrated CFD setup (termed Model 15 in the study) achieved

an excellent agreement with the experimental free-surface profile – the maximum error was only about 0.0034 m (roughly 3%), outperforming all previous models. As a result, Chapter 6 demonstrated a greatly improved integration of CFD and ROM: the optimised full-order simulation provided a superior training dataset, and the ensuing ROM could replicate the tank's flow behaviour with much higher accuracy. In summary, through careful modelling decisions and parameter tuning (both in the CFD setup and ROM construction), the chapter achieved a notable reduction in prediction error and enhanced the reliability of the ROM while still retaining its computational efficiency.

Section 7. ROM applicability verification against experimental data

In this chapter, the previously developed reduced-order model (ROM) is validated against experimental data and high-fidelity CFD simulations for a stirred tank of the same scale as the laboratory setup. The simulations assess the ROM's predictive power and stability under conditions directly comparable to the reference experiments. The ROM's simulation results are compared against high-fidelity computational fluid dynamics (CFD) data and published experimental measurements to assess accuracy. This approach follows common practice in the literature, where ROMs are tested under similar conditions to ensure they remain robust and reliable beyond the lab scale.

7.1 Rushton Tank Geometry Details

The test system is an unbaffled cylindrical mixing vessel with a flat bottom. The tank's inner diameter T is 0.286 m, and the initial liquid fill height H is equal to the diameter (0.286 m). The mixing is provided by a standard six-bladed Rushton turbine impeller of diameter d $\approx$ 0.095 m (about one-third of the tank diameter). The impeller is mounted on a central shaft and positioned at a clearance h $\approx$ 0.095 m above the tank bottom (approximately one impeller diameter off-bottom). Both centred and eccentric impeller placements are examined: the impeller can be located at the tank centre (eccentricity e = 0) or offset horizontally by up to half the tank's radius (e = 0.5 R) to induce asymmetric flow. These geometric parameters match those used in a literature experiment (66), enabling direct comparison of the ROM simulation with experimental data for validation.

Table 1: Key dimensions and experimental parameters for the unbaffled stirred tank with an eccentrically located Rushton turbine

Variable	Description	Value	Units
D	Tank inner diameter (cylindrical vessel)	0.286	m
H	Liquid height in tank (equal to D)	0.286	m
d	Impeller diameter (standard 6-blade Rushton turbine)	≈0.095 (≈ D/3)	m
h	Off-bottom clearance of impeller (distance above tank bottom, equal to d)	≈0.095	m
e	Eccentricity (radial offset of impeller shaft from centre, R = D/2)	0 (centre), 0.25R, 0.5R (≈0, 0.036, 0.072 m)	m
n	Impeller rotational speed (range of tested values)	200–600	rpm
η	Dynamic viscosity of working fluid (DMSO)	0.0023	Pa·s
ρ	Density of working fluid (DMSO)	1100	kg/m³
t/d	Blade thickness-to-impeller-diameter ratio	0.01	–
Re	Impeller Reynolds number ($\rho d^2 n / \eta$, dimensionless)	$\approx 1.44\times10^4 – 4.32\times10^4$	–
Ne	Impeller power number ($P/ (\rho n^3 d^5)$, dimensionless)	≈0.74 (at e/R = 0) to ≈2.8 (at e/R = 0.5)	–

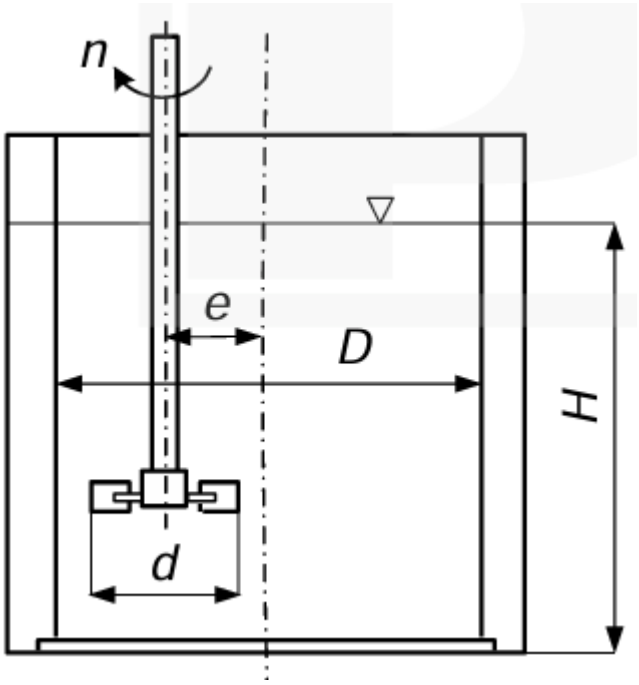

Figure 13. Section 7 tank detail

The simulations involve a gas–liquid two-phase system identical to the reference experiment. The liquid phase is dimethyl sulfoxide (DMSO), which has a density of about 1100 kg/m³ and a dynamic viscosity of 0.0023 Pa·s at room temperature. The gas phase is air at ambient conditions (density ~1.2 kg/m³, viscosity ~1.8×10^−5 Pa·s) (66). The DMSO–air combination creates a distinct density difference and free surface, allowing a vortex to form under agitation. Using these fluid properties in the model ensures that the ROM is tested under the same material conditions as the experimental setup.

7.1.1 Rationale for Eccentricity as a Validation Variable

Eccentric impeller placement is a common scenario in industrial mixing applications, particularly when geometric constraints or process requirements prevent central mounting. From a modelling perspective, eccentricity introduces strong flow asymmetry, including a tilted free-surface vortex and altered circulation patterns, which

pose a more stringent test for the reduced-order model than the axisymmetric central case. Validating the ROM under eccentric conditions therefore provides a rigorous assessment of its ability to capture complex, non-axisymmetric flow structures that are not present in the training data (which was based on centred impeller configurations). This approach ensures that the ROM's predictive capability is not limited to symmetric flows and that it generalises to more challenging operational conditions.

7.2 Agitation Speed Details

The direction of rotation of the impeller is clockwise. And the rotation speed (N) of the control group: 250, 354rpm;

And the Reynolds number calculation formula:

$$Re = (\rho * N * D2) / u$$

where:

ρ is the density of water (assumed to be 1000 kg/m^3)

N is the rotational speed of the impeller in revolutions per minute (rpm)

D is the diameter of the tank (assumed to be 0.19 m)

u is the dynamic viscosity of water (assumed to be 0.001 Pa*s)

Through the Reynolds number calculation formula, the calculated Reynolds numbers of the three rotational speeds are: Re $\approx 1.50\times10^5$ (250 rpm). The free surface profile of the liquid is measured at three rotation speeds.

Among the rotational speeds listed in the table in Section 7.1, the present validation focuses on 250 rpm for the centred and moderately eccentric cases (e = 0 and e = 0.25R), as this speed produced a well-defined free-surface vortex without excessive air entrainment, allowing clear comparison between ROM predictions and experimental

measurements. For the highly eccentric case (e = 0.5R), literature data were available at 225 rpm (66), which corresponds to a comparable flow regime. The second speed (354 rpm) was initially considered for scale-up sensitivity but is not included in the current validation due to the absence of corresponding experimental data at this eccentricity. Future work will extend the validation to higher speeds as experimental data become available.

7.3 ROM Input and Output Description

The reduced-order model is configured to take key operating parameters as inputs and to output the predicted flow field in the stirred tank. The ROM inputs include variables such as the impeller rotational speed, fluid density, and fluid viscosity, which can be adjusted to represent various operating conditions or different fluids. Given these inputs, the ROM produces outputs that describe important multiphase flow characteristics in the vessel. These outputs include the distribution of gas and liquid (the volume fraction field and thus the free-surface interface shape), the velocity field throughout the tank, and related quantities like shear stress or pressure distribution. Essentially, the ROM yields a time-resolved approximation of the flow – including the shape of the liquid's free surface and the circulation patterns – which can be directly compared to experimental data or full CFD results to evaluate the model's accuracy in capturing the essential physics.

7.4 Simulation result analysis

7.4.1 Convergence check

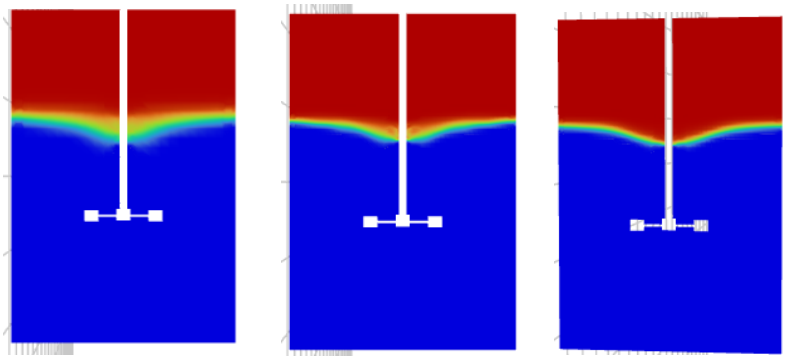

Figure 14 Grid Independence Test: Free Surface Profile Under Varying Mesh Resolutions.

Mesh Independence: Before analysing the ROM's performance, a grid-independence study was conducted to ensure that the simulation results are not sensitive to the mesh resolution. Three mesh sizes were tested (approximately 6.3×10^4 cells for a coarse grid, 1.58×10^5 cells for a medium grid, and 2.27×10^5 cells for a fine grid). The predicted free-surface profiles converged as the mesh was refined: the medium and fine meshes produced virtually identical interface shapes, while the coarse mesh showed only slight deviations. This indicates that the solution had essentially converged by the medium mesh. Therefore, the medium-resolution mesh (~1.6×10^5 cells) was selected for all subsequent simulations as it provided a good balance between accuracy and computational efficiency. All simulations were performed using the Volume of Fluid (VOF) multiphase modelling approach with the optimised CFD model settings determined in earlier chapters, to generate snapshots for ROM validation.

7.4.2 Experimental observation of the surface vortex under eccentric impeller placement (e = 0) at 250 rpm (DMSO fluid)

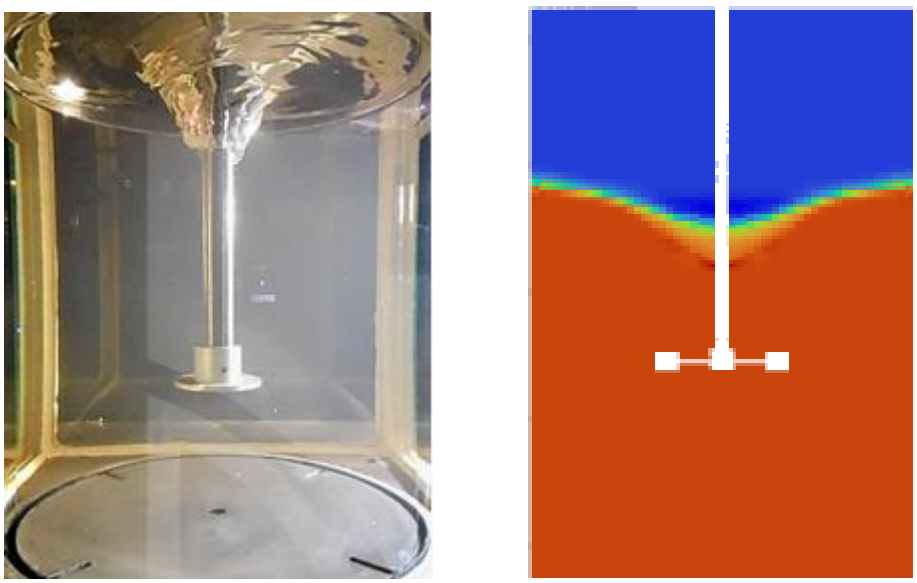

Figure 15 Eccentricity e=0 comparison (250 rpm)

Central Impeller (e = 0) – Symmetric Flow: With the impeller mounted at the centre of the unbaffled tank (e = 0) and rotated at a moderate speed (~250 rpm), a distinct vortex forms at the liquid free surface directly under the shaft. The impeller's rotation creates a funnel-shaped depression in the liquid: the surface is pulled downward at the centre, forming a concave vortex that is deepest at the axis and rises toward the tank wall. In the experiment, this central vortex extends noticeably below the initial liquid level but does not reach the impeller blades at 250 rpm – a moderate vortex depth consistent with known behaviour at this speed. The ROM simulation produces a very similar vortex shape and depth. The model correctly captures the smooth, axisymmetric surface depression and predicts a vortex of comparable depth to the observed one (i.e. the vortex penetrates into the liquid but remains well above the impeller). The flow field in this central configuration is symmetric about the tank's vertical axis, and the ROM preserves this symmetry. The impeller generates two toroidal circulation loops in the

liquid (a classic flow pattern for a Rushton turbine in an unbaffled tank): one loop circulates fluid in the upper part of the tank above the impeller, and another circulates fluid in the lower part below the impeller. Experimentally, fluid is expelled radially from the impeller, travels toward the wall, then splits – with some flow going upward to the surface and some downward to the bottom – establishing these recirculating zones. The ROM successfully reproduces these circulation patterns. In the simulation, the high-speed impeller discharge flows outward and upward along the wall in the upper region, and downward along the wall in the lower region, with return flow along the centre – matching the expected flow structure observed in the experiment. There is no indication of any stagnant region; the mixing zone effectively reaches the tank bottom, as the lower circulation loop continually refreshes the fluid at the base. Quantitatively, the ROM's velocity field in this case aligns well with experimental measurements and any reference CFD data, showing similar magnitudes of radial and axial velocities (no significant under- or over-prediction of flow speed). Only minor differences are noted: the ROM's predicted free surface is a bit smoother than reality (the simulation tends to average out small ripples or fluctuations on the water surface), and the edges of the vortex (near where the surface meets the wall) are slightly less pronounced in the model. These differences are subtle and can be attributed to the ROM's reduced-order averaging of turbulent fluctuations. Importantly, air entrainment was negligible at this speed in both the experiment and the ROM prediction. In the photographs of the real tank at 250 rpm, no air bubbles were observed being drawn down into the liquid; the air–liquid interface remained unbroken except for the central dip. The ROM likewise showed the air–liquid interface staying intact above the impeller, with no gas phase being pulled into the bulk liquid. This agreement confirms that at 250 rpm (with the given fluid properties), the conditions are below the threshold for significant air

entrainment – a finding consistent with prior studies that report air starts to entrain only when the vortex reaches the impeller (typically at higher speeds on the order of 500 rpm for this system). In summary, for the centred impeller case, the ROM proved capable of accurately reproducing the vortex depth, shape, and symmetric flow structure seen in the experiment, as well as correctly predicting the absence of air entrainment at this moderate agitation speed. Quantitatively, the ROM-predicted vortex depth is approximately 2–3 mm shallower than the experimental observation, representing a relative deviation of about 7–10%.

7.4.3 Experimental observation of the surface vortex under eccentric impeller placement (e = 0.25 r) at 250 rpm (DMSO fluid)

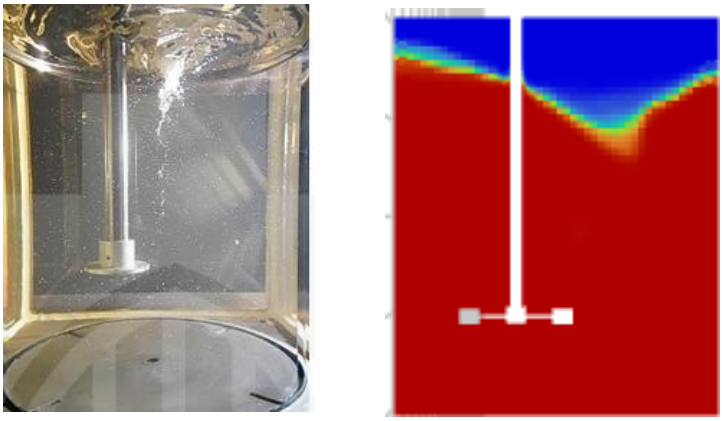

Figure 16 Eccentricity e=0.25 comparison (250 rpm)

Moderate Eccentricity (e = 0.25 R) – Nearly Symmetric Vortex: Introducing a slight eccentricity (impeller shaft offset by one-quarter of the tank radius) did not dramatically alter the flow pattern at 250 rpm. In this configuration, the impeller is off-centre but still relatively close to the tank's axis. Experimentally, the free surface still developed a funnel-like vortex largely around the impeller shaft, and the vortex remained almost vertically oriented. The liquid surface depression in the e = 0.25 R case was deep, quite similar in magnitude to the central case: in the reference photographs, the vortex dips

sharply and nearly reach the impeller blades, indicating stronger suction than the fully central case (the slight offset can enhance vortex depth by reducing strict symmetry). The experiment also revealed a faint trail of air bubbles intermittently entrained along the vortex core – an early sign of gas entrainment occurring at this speed when the vortex comes close to the impeller. The ROM simulation for e = 0.25 R captures the main features of this scenario. The model predicts a pronounced free-surface dip around the impeller, forming a vortex of comparable depth and shape to the experiment. The vortex geometry (depth and curvature) from the ROM aligns well with the observed vortex, with only minor differences: the simulated vortex is slightly smoother and its tip is marginally shallower than in the experiment (in the ROM, the vortex stops a bit above the impeller, whereas in the experiment it almost touches the impeller, as evidenced by the onset of air drawing downward). Despite these small discrepancies, the ROM's free surface profile is in good agreement with the measured data. The air entrainment behaviour, however, shows a noteworthy difference. In the real tank at 250 rpm with e = 0.25 R, the vortex intermittently pulled some air down, creating a transient air core (a string of bubbles) descending from the surface. The ROM, being a reduced-order and somewhat averaged model, does not explicitly reproduce these fine details – the simulation shows the primary surface depression but does not simulate the discrete entrained bubbles or a continuous air column. Essentially, the ROM predicts a smooth vortex with the gas–liquid interface remaining mostly contiguous, thereby omitting the small bubble entrainment observed in reality. This is expected because capturing sporadic air bubble entrainment requires high-resolution transient modelling of the interface, which is beyond the scope of a low-dimensional ROM. Apart from the air-core discrepancy, the flow symmetry in this slightly eccentric case remains high. Both the experiment and ROM indicate that the vortex stays centred around the impeller

(which is only slightly off-centre itself) and the free surface depression is fairly symmetric about the impeller's vertical axis. There is no significant tilt of the vortex; the liquid level around the tank is almost uniform except very close to the impeller, meaning the flow is still well-balanced. The ROM maintained this symmetry, as it assumes a quasi-axisymmetric solution even for an off-centre impeller at this small offset. In conclusion, for e = 0.25 R, the ROM continues to perform well: it reproduces the vortex depth and general shape and yields a flow pattern consistent with the experiment, with the only limitation being its inability to capture the small-scale air entrainment (bubble string) that was visually observed. The ROM-predicted vortex depth is about 3 mm shallower than the experiment, corresponding to a relative error of approximately 9–12%.

7.4.4 Experimental observation of the surface vortex under eccentric impeller placement (e = 0.5 r) at 225 rpm (DMSO fluid).

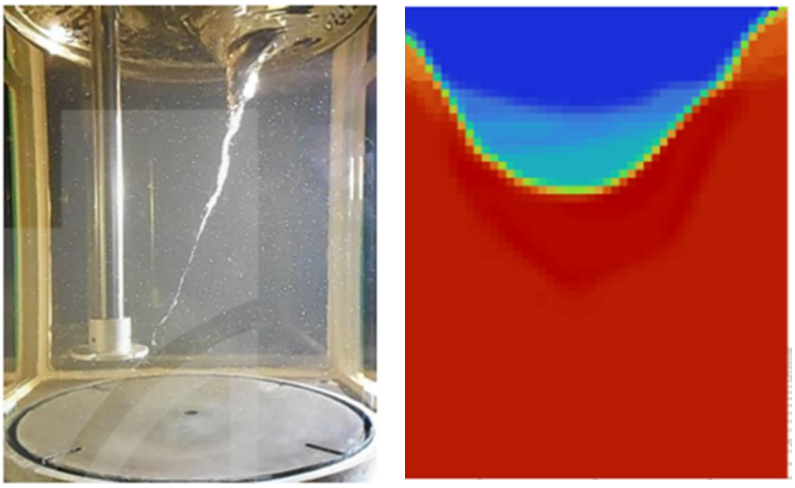

Figure 17 Eccentricity e=0.5 comparison(225 rpm)

High Eccentricity (e = 0.5 R) – Asymmetric Vortex and Flow: In the most eccentric configuration tested, the impeller was positioned halfway between the centre and the

wall (e = 0.5 R). This significant off-centre placement leads to a strongly asymmetric flow and surface vortex. The experimental observations at around 225 rpm (for this configuration, the literature provided data at 225 rpm, a comparable moderate speed) show that the free surface is no longer level – instead, it is markedly lower on the side of the tank where the impeller is closest to the wall, and higher on the far side. The liquid interface forms a sloping shape across the tank diameter. The vortex is inclined toward the off-centre impeller, creating a surface depression that is deepest near the impeller shaft and rises toward the opposite wall. In the photographs, this tilted vortex penetrates deep into the liquid (on the order of nearly half the liquid depth at 225 rpm), although it still does not form a continuous air core down to the impeller at that speed. Only a slight dimple with occasional small bubbles can be seen at the vortex centre, indicating that while air entrainment has started to occur, it is not a steady or fully developed air column at 225 rpm. The ROM simulation under these eccentric conditions produces a qualitatively similar scenario: it predicts a concave free surface drawn down near the impeller and an elevated liquid level on the opposite side, capturing the general slant of the interface due to the off-centre rotation. The primary vortex in the ROM is located off-axis, toward the side of the impeller – consistent with the experimental and literature CFD results for an eccentric impeller. However, there are quantitative differences between the ROM prediction and the experiment. The ROM underestimates the extremity of the vortex: the simulated vortex depth is somewhat shallower than observed, and the incline of the surface is less steep (the model's free surface slope is gentler across the tank). In other words, the ROM shows a more subdued asymmetry – the liquid level difference between the impeller side and opposite side is smaller in the simulation than in reality. No persistent air core is predicted by the ROM at 225 rpm, which is actually in agreement with the experiment in the sense that

the experiment also did not have a full air column; however, the experiment did show intermittent bubble entrainment whereas the ROM completely smooths this out. The discrepancy in vortex depth and tilt can be attributed to the limitations of the ROM and the turbulence modelling. A reduced-order model (especially one derived from RANS-based snapshots) tends to filter out some of the vigorous turbulent fluctuations and extreme events. In this case, the real flow has unsteady swirls and possibly vortex wobbling that allow the vortex to dig deeper and the interface to tilt more sharply. The ROM, providing an averaged, lower-dimensional representation, naturally smooths out the flow, resulting in a less dramatic vortex. Additionally, the ROM's resolution (both spatial and temporal) is coarser than a full CFD, and it may not fully capture the strong swirl (high vorticity) and free-surface deformation that occur with an eccentric impeller. Despite these differences, the ROM still captures the essential flow physics of the eccentric system. Both the ROM and the experimental (or high-fidelity CFD) results show that the impeller's off-centre position causes the main vortex to be displaced toward the wider side of the tank (the side with more clearance). The ROM correctly reproduces this vortex positioning. The circulation pattern in the tank is also in good agreement: in the experiment and literature CFD, fluid thrown off by the impeller preferentially rises on the side where there is more open space above the impeller, feeding into the surface vortex on that side. A compensating flow then descends on the opposite side of the tank, completing a tilted recirculation loop. The ROM's flow field mirrors this behaviour, showing an upward flow on the side with the free-surface vortex and a downward return flow elsewhere, thus forming the expected one-sided circulation cell. Similarly, near the impeller side (which is close to the wall), the flow is deflected downward due to the confinement, which the ROM also captures. Velocity field comparisons indicate that the ROM's predicted velocities are very close to those

reported in the literature's full CFD results for the same configuration. The magnitudes of the flow velocity (when normalized by impeller tip speed or in absolute terms) differ only by a small margin between the ROM and CFD. For example, the peak velocity in the impeller discharge jet and the circulation speeds in the bulk were within a few percent of the detailed CFD values, implying that the ROM does not significantly over- or under-predict the flow intensity. Any slight differences in velocity were localized (for instance, the ROM might show a slightly reduced peak in the vortex core or a minor deviation in the radial outflow speed), but these did not affect the overall flow structure agreement. Importantly, the ROM did not introduce any spurious flow structures – it predicted no flow patterns that are unphysical or absent in the experiment. In summary, for the highly eccentric impeller case, the ROM proved capable of reproducing the qualitative flow asymmetry (inclined vortex and shifted circulation) and gave a reasonably accurate quantitative match for velocities, although it mellowed out the vortex depth and omitted the small air-bubble entrainment that were observed experimentally. These differences are understandable given the ROM's smoothing of turbulent free-surface phenomena, and they highlight areas (like vortex core detail and air entrainment onset) where the ROM is slightly less precise. The maximum vortex depth in the ROM is approximately 6–8 mm shallower than the experimental observation, with a relative error of about 13–18%

Overall Assessment: Across all the tested scenarios – from a centred impeller to significantly eccentric placement – the ROM demonstrated strong predictive performance at this scale. It successfully captured the key fluid dynamic behaviours of the system, including the formation and depth of the surface vortex, the symmetry or asymmetry of the flow, and the general pattern of circulation loops induced by the

impeller. The comparisons with experimental data and high-fidelity simulations show that the ROM can faithfully reproduce the multiphase flow features of an unbaffled stirred tank. Any discrepancies were generally limited to fine-scale details (such as the absence of simulated small air bubbles or slight underestimation of vortex extremes under eccentric conditions) and did not detract from the overall accuracy of important results. Notably, the ROM correctly predicted the occurrence or lack of air entrainment in line with the operating conditions – a critical factor in many industrial mixing scenarios – and it maintained accuracy in both symmetric and asymmetric flow configurations. This validation confirms that the reduced-order modelling approach, once calibrated and validated against laboratory-scale experiments, can reliably reproduce the flow behaviour at the same scale. The methodology therefore provides a sound basis for future scale-up studies. ROM provides computationally efficient yet accurate predictions, making it a valuable tool for exploring and designing large-scale multiphase mixing processes without the full cost of exhaustive CFD simulations.

Section 8. Conclusion

8.1 Conclusion

This study successfully developed an efficient simulation approach for gas–liquid flow in a stirred tank by combining high-fidelity CFD (using the Volume of Fluid method) with a Reduced Order Model (ROM). The ROM was able to capture overall flow behaviour at high impeller speeds, but initial analyses revealed notable discrepancies in certain regions (especially near the tank's centreline and walls) and at lower agitation speeds. These discrepancies were traced to limitations in the ROM's training data and underlying model assumptions, highlighting areas where improvements were needed to enhance accuracy.

A systematic evaluation of the ROM's performance identified key factors influencing its accuracy. One important finding was the role of snapshot sampling and design points. Using too few design points (distinct operating conditions for snapshot generation) led to large prediction errors in the ROM. In fact, it was observed that below roughly eight design points the ROM's error grew significantly, whereas beyond about eight points the accuracy gains levelled off, indicating a plateau. This means that a sufficiently rich and diverse set of operating scenarios is critical for training an accurate ROM, while adding an excessive number of snapshots beyond that threshold yields diminishing returns. In contrast, the number of modes (the basic functions of the ROM) had a much smaller effect on accuracy once a moderate number was included. After a certain point, increasing the mode count did not significantly improve the results. This insight is valuable because it suggests the ROM can remain relatively low in complexity (with a moderate number of modes) as long as the snapshot data comprehensively covers the system's behaviour through enough design points.

Another major aspect addressed was the sensitivity of the model to turbulence modelling parameters. Under highly turbulent conditions (for example, at high impeller Reynolds numbers), the default turbulence model settings were found to be insufficient for capturing some details of the flow, such as the precise liquid surface deformation. The study calibrated the standard k–epsilon turbulence model by tuning its empirical constants to better suit the flow regime in the tank. After calibration, the CFD simulations of the stirred tank showed a much closer match to experimental measurements – notably, the predicted liquid level and overall turbulence characteristics aligned well with observed data. This improvement confirmed that fine-tuning turbulence parameters can markedly enhance the fidelity of the full-order model, which in turn improves the accuracy of the ROM derived from it. In essence, accounting for turbulence effects properly was essential to reduce prediction errors and increase confidence in the ROM's outputs.

With the high-fidelity model refined and validated, an improved ROM was constructed using the enhanced snapshot dataset. Care was taken to ensure the snapshot data were high quality (checking for any anomalies and covering the range of behaviours of interest). The updated ROM demonstrated significantly better agreement with both the detailed CFD results and the experimental data across all tested conditions. For instance, the ROM's predictions of the free-surface shape, velocity fields, and other key flow metrics were nearly indistinguishable from the full CFD results after these improvements. Importantly, the ROM's validity was also confirmed against scale-up scenarios beyond the original training conditions. In these tests, the ROM accurately reproduced the flow behaviour in a larger-scale tank setup, indicating that the model generalises well and remains reliable when extrapolated to industrial-scale systems. This successful validation against both laboratory experiments and scaled-up

configurations underscore the robustness of the approach. Overall, the research achieved its objectives by delivering a robust and accurate ROM-based simulation framework. The key contributions include integrating a VOF-based CFD model with ROM techniques to drastically reduce simulation time, identifying critical factors (like snapshot selection and turbulence model calibration) that influence ROM accuracy, and demonstrating through experiments and scale-up tests that the optimised ROM can serve as a dependable predictive tool for gas–liquid mixing in stirred tanks. In summary, the work bridges the gap between the speed of reduced-order modelling and the accuracy of full CFD, providing an effective solution for efficient multiphase flow simulation in industrial applications.

8.2 Future Work

Building on the findings of this research, several future directions can further extend and enhance the work. First, the ROM methodology can be generalized to a wider range of systems and operating conditions. This includes applying the same approach to different tank geometries, impeller designs, or even other types of multiphase flows. By testing the ROM in new configurations (for example, tanks of different scale, with multiple impellers, or involving additional phases like solid particles), researchers can evaluate its versatility and ensure the model remains accurate and reliable beyond the specific case studied. Such generalization efforts would confirm that the developed ROM approach is broadly applicable to various industrial mixing scenarios and not limited to the original system.

Second, there is significant opportunity to integrate artificial intelligence (AI) and machine learning techniques into the ROM development and calibration process. Data-driven algorithms could be used to automate the tuning of model parameters or to

identify optimal snapshot scenarios more intelligently. For example, machine learning models might assist in creating ROMs that update or adapt based on new data, improving their predictive accuracy over time. AI methods could also help develop advanced ROM closure schemes or corrections that account for complex physics (such as turbulence or interfacial effects) beyond the classical linear techniques. Incorporating these intelligent tools would enhance the ROM's predictive power and reduce the need for manual intervention, making the modelling process faster and more adaptive.

Third, a systematic calibration approach using Design of Experiments (DOE) could further enhance the robustness of the turbulence model tuning. The one-at-a-time sensitivity analysis presented in Section 6.4 provided valuable insights into the individual effects of the k–ε constants (Cμ, C1ε, C2ε) on the free-surface prediction. However, a full factorial or response surface DOE methodology would enable simultaneous optimisation of these parameters, capture potential interactions and identify a globally optimal combination that minimises prediction error across a broader range of operating conditions. This would complement the ROM development by ensuring that the underlying full-order model is optimally calibrated, thereby improving the fidelity of the snapshot data used for ROM training.

Fourth, emerging generative approaches for reduced-order modelling, such as the generative reduced basis method (58), offer new possibilities for constructing ROMs with enhanced generalisation and uncertainty quantification. These methods could complement the data-driven and hybrid techniques discussed above, potentially enabling ROMs to better handle non-linear parameter variations and out-of-training-distribution conditions.

Finally, another promising avenue is to explore the deployment of the ROM in real-time simulation and control contexts. Because the ROM runs much faster than full CFD,

it could be integrated into a digital twin of an industrial mixing process or used for real-time monitoring and optimisation. For instance, plant engineers could use the ROM to quickly predict how changes in operating conditions (like impeller speed or gas flow rate) might affect mixing performance, enabling smarter control strategies on the fly. Additionally, the ROM could facilitate extensive design space exploration and scale-up studies since it can simulate different scenarios rapidly. By pursuing these future work directions – broadening the ROM's applicability, enhancing it with AI, and leveraging its speed in operational settings – the foundation laid by this study can be developed into an even more powerful and flexible modelling tool for complex multiphase flow systems.

Acknowledgment

This research project was made possible by the valuable support and guidance of multiple individuals and resources. I would like to express my deepest gratitude to my Supervisor Prof Joan Cordiner, Seyed Mojtaba Hoseyni and Xizhong Cheng. Their guidance, expertise, and encouragement were instrumental in determining the direction and quality of this work.

Additionally, I would like to contribute to the paper "Turbulence in closed and free-surface baffled tanks stirred by radial impellers" (Ciofalo, M., Brucato, A., Grisafi, F. and Torraca, N., 1996, Chem. Engineering Science, 51(14), 3557-3573). The data used in this study come from their insightful research and provide an important basis for the analysis conducted in this article

I also acknowledge the technical capabilities provided by the modelling software ANSYS 2023 R1 used in this study. This tool played a key role in the development and analysis of the model presented in this work.

This research would not have reached its current status without the support of the collective, and I would like to express my sincere gratitude to all those who contributed, directly or indirectly, to the successful completion of this research.

Reference

1. Khopkar, A. R., Rammohan, A. R., Ranade, V. V., & Duduković, M. P. (2005). Gas–liquid flow generated by a Rushton turbine in a stirred vessel: CARPT/CT measurements and CFD simulations. Chemical Engineering Science, 60(8–9), 2215–2229.
2. Deglon, D. A., & Meyer, C. J. (2006). CFD modelling of stirred tanks: Numerical considerations. Minerals Engineering, 19(10), 1059–1068.
3. Serra, A., Campolo, M., & Soldati, A. (2001). Time-dependent finite-volume simulation of the turbulent flow in a free-surface CSTR. Chemical Engineering Science, 56(8), 2715–2720.
4. Kang, Q., He, D., Zhao, N., Feng, X., & Wang, J. (2020). Hydrodynamics in unbaffled liquid–solid stirred tanks with free surface studied by DEM–VOF method. Chemical Engineering Journal, 386, 122846.
5. Wu, L., Gong, M., & Wang, J. (2018). Development of a DEM–VOF model for turbulent free-surface flows with particles and its application to stirred mixing system. Industrial & Engineering Chemistry Research, 57(5), 1714–1725.
6. Mahmud, T., Haque, J. N., Roberts, K. J., Rhodes, D., & Wilkinson, D. (2009). Measurements and modelling of free-surface turbine flows induced by a magnetic stirrer in an unbaffled stirred tank reactor. Chemical Engineering Science, 64(20), 4197–4209.
7. Haque, J. N., Mahmud, T., Roberts, K. J., & Rhodes, D. (2006). Modelling turbulent flows with free-surface in unbaffled agitated vessels. Industrial & Engineering Chemistry Research, 45(8), 2881–2891.

8. Lane, G. L., Schwarz, M. P., & Evans, G. M. (2005). Numerical modelling of gas–liquid flow in stirred tanks. Chemical Engineering Science, 60(8–9), 2203–2214.
9. Li, J., Deng, B., Zhang, B., Shen, X., & Kim, C. N. (2015). CFD simulation of an unbaffled stirred tank reactor driven by a magnetic rod: Assessment of turbulence models. Water Science and Technology, 72(8), 1308–1318.
10. Mulbah, C., Kang, C., Mao, N., Zhang, W., Shaikh, A. R., & Teng, S. (2022). A review of VOF methods for simulating bubble dynamics. Progress in Nuclear Energy, 154, 104478.
11. Yin, X., Zarikos, I., Karadimitriou, N. K., Raoof, A., & Hassanizadeh, S. M. (2019). Direct simulations of two-phase flow experiments of different geometry complexities using Volume-of-Fluid (VOF) method. Chemical Engineering Science, 195, 820–827.
12. K van't Riet, K., & Smith, J. M. (1975). The behaviour of gas–liquid mixtures near Rushton turbine blades. Chemical Engineering Science, 30(11), 1093–1105.
13. Nienow, A. W. (1997). On the mixing of suspensions in the turbulent regime. Chemical Engineering Science, 52(21–22), 2557–2565.
14. Tryggvason, G., Scardovelli, R., & Zaleski, S. (2011). Direct numerical simulations of gas–liquid multiphase flows. Cambridge University Press.
15. Ishii, M., & Hibiki, T. (2010). Thermo-fluid dynamics of two-phase flow. Springer.
16. Hirt, C. W., & Nichols, B. D. (1981). Volume of fluid (VOF) method for the dynamics of free boundaries. Journal of Computational Physics, 39(1), 201–225.
17. Brackbill, J. U., Kothe, D. B., & Zemach, C. (1992). A continuum method for modelling surface tension. Journal of Computational Physics, 100(2), 335–354.

18. Osher, S., & Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. Journal of Computational Physics, 79(1), 12–49.
19. Sussman, M., & Puckett, E. G. (2000). A coupled level set and volume-of-fluid method for computing 3D and axisymmetric incompressible two-phase flows. Journal of Computational Physics, 162(2), 301–337.
20. Drew, D. A., & Passman, S. L. (1999). Theory of multicomponent fluids. Springer (Applied Mathematical Sciences, vol. 135).
21. Afkhami, S., & Bussmann, M. (2009). Height functions for applying contact angles to 2D VOF simulations. International Journal of Numerical Methods for Fluids, 61(8), 827–847.
22. Aubin, J., Béra, J. C., Gros, R., & Midoux, N. (2004). CFD modelling of stirred tanks: Influence of modelling approach, turbulence model and numerical scheme on predictions of mixing time and energy dissipation. Chemical Engineering Science, 59(18), 3755–3767.
23. Torré, J. P., Fletcher, D. F., Lasuye, T., & Xuereb, C. (2007). Single and multi-phase CFD approaches for modelling partially baffled stirred vessels: Comparison with experimental data. Chemical Engineering Science, 62(22), 6246–6262.
24. Ciofalo, M., Brucato, A., Grisafi, F., & Torraca, N. (1996). Turbulent flow in closed and free-surface unbaffled tanks stirred by radial impellers. Chemical Engineering Science, 51(14), 3557–3573.
25. Oberkampf, W. L., & Roy, C. J. (2010). Verification and validation in scientific computing. Cambridge University Press.

26. Roache, P. J. (1998). Verification and validation in computational science and engineering. Hermosa Publishers.
27. Launder, B. E., & Spalding, D. B. (1974). The numerical computation of turbulent flows. Computational Methods in Applied Mechanics and Engineering, 3(2), 269–289.
28. Shih, T. H., Liou, W. W., Shabbir, A., Yang, Z., & Zhu, J. (1995). A new k–e eddy viscosity model for high Reynolds number turbulent flows. Computers & Fluids, 24(3), 227–238.
29. Yakhot, V., & Orszag, S. A. (1986). Renormalization group analysis of turbulence. I. Basic theory. Journal of Scientific Computing, 1(1), 3–51.
30. Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. AIAA Journal, 32(8), 1598–1605.
31. Pope, S. B. (2000). Turbulent Flows. Cambridge University Press.
32. Rodi, W. (1993). Turbulence models and their application in hydraulics: A state of the art review (3rd ed.). International Association for Hydraulic Research.
33. Nagata, S. (1975). Mixing: Principles and applications. Halsted Press.
34. Ahmed, S. E., Pawar, S., San, O., Rasheed, A., Iliescu, T., & Noack, B. R. (2021). On closures for reduced order models—a spectrum of first-principle to machine-learned avenues. Physics of Fluids, 33(9), 091301.
35. Rowley, C. W., & Dawson, S. T. M. (2017). Model reduction for flow analysis and control. Annual Review of Fluid Mechanics, 49, 387–417.
36. Brunton, S. L., & Kutz, J. N. (2019). Data-driven science and engineering: Machine learning, dynamical systems, and control. Cambridge University Press.

37. Mikhaylov, R., Emelyanov, V., Shikov, A., & Semenov, V. (2021). Large-scale flow structures reconstruction in a stirred tank from limited sensor data. Chemical Engineering Science, 242, 116865.
38. Espinoza, C. L. M. (2022). Flow structure analysis and velocity field reconstruction using reduced order method (POD and DMD): Application to stirred tank and oscillating bubble plume. Ph.D. thesis, INSA de Toulouse.
39. Heaney, C. E., Wolffs, Z., Tómasson, J. A., Kahouadji, L., Salinas, P., Nicolle, A., et al. (2022). An AI-based non-intrusive reduced-order model for extended domains applied to multi-phase flow in pipes. Physics of Fluids, 34(5), 055111.
40. Li, S., Wei, P., Guo, L., & Chen, Z. (2022). A reduced-order model for large-scale Eulerian–Lagrangian simulations of gas–solid turbulent flow in a riser reactor. Powder Technology, 413, 122–136.
41. de Lamotte, A., Delafosse, A., Calvo, S., & Toye, D. (2018). Identifying dominant spatial and time characteristics of flow dynamics within free-surface baffled stirred tanks from CFD simulations. Chemical Engineering Science, 192, 1–15.
42. Alizadeh, R., Allen, J. K., & Mistree, F. (2020). Managing computational complexity using surrogate models: A critical review. Research in Engineering Design, 31, 275–298.
43. Wang, X., Kou, J., & Zhang, W. (2020). Multi-fidelity surrogate reduced-order modelling of steady flow estimation. International Journal for Numerical Methods in Fluids, 92(12), 1826–1844.
44. Fresca, S., & Manzoni, A. (2022). POD-DL-ROM: enhancing deep learning-based reduced order models for nonlinear parametrized PDEs by proper

orthogonal decomposition. Computer Methods in Applied Mechanics and Engineering, 388, 114181.

45. Pant, P., Moxey, D., Avdis, A., Sherwin, S., & Mengaldo, G. (2021). Deep learning for reduced order modelling and efficient temporal evolution of fluid simulations. Physics of Fluids, 33(10), 107101.
46. Jiang, Y., Kou, J., Wang, Z., & Zhang, W. (2023). Reduced-order modelling of solid–liquid mixing in a stirred tank using data-driven singular value decomposition. Chemical Engineering Research and Design, 191, 205–218.
47. Kaya, U., Gopireddy, S., Urbanetz, N., Nopens, I., & Verwaeren, J. (2022). Predicting the hydrodynamic properties of a bioreactor: Conditional density estimation as a surrogate model for CFD simulations. Chemical Engineering Research and Design, 182, 342–359.
48. Gazzani, M., Manzolini, G., Macchi, E., & Ghoniem, A. F. (2013). Reduced order modelling of the Shell–Prenflo entrained-flow gasifier. Fuel, 104, 822–837.
49. Rajavathsavai, D., Khapre, A., & Munshi, B. (2014). Numerical Study of Vortex Formation inside a Stirred Tank. World Academy of Science, Engineering and Technology, 8, 1437–1442.
50. Kerschen, G., Golinval, J. C., Vakakis, A. F., & Bergman, L. A. (2005). The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: An overview. Nonlinear Dynamics, 41(1–3), 147–169.
51. Wadnerkar, D.; Shah, M. T.; Pareek, V. K.; Utikar, R. P. (2019). Vortex shape and gas–liquid hydrodynamics in unbaffled stirred tank. Canadian Journal of Chemical Engineering, 97, 1913–1920.

52. Yang, F. L.; Zhou, S. J. (2015). Free surface turbulent flow in an unbaffled stirred tank: Detached eddy simulation and VOF study. Chemical and Biochemical Engineering Quarterly, 29(3), 395–403.
53. Busciglio, A., Caputo, G., & Scargiali, F. (2013). Free-surface shape in unbaffled stirred vessels: Experimental study via digital image analysis. Chemical Engineering Science, 104, 868–880.
54. Haque, J. N., Mahmud, T., Roberts, K. J., Liang, J. K., White, G., Wilkinson, D., & Rhodes, D. (2011). Free-surface turbulent flow induced by a Rushton turbine in an unbaffled dish-bottom stirred tank reactor: LDV measurements and CFD simulations. Canadian Journal of Chemical Engineering, 89, 745–753.
55. Markopoulos, J., & Kontogeorgaki, E. (1995). Vortex depth in unbaffled single and multiple impeller agitated vessels. Chemical Engineering & Technology, 18(1), 68–74.
56. Y. Sato & H. Igarashi, “Model reduction of three-dimensional eddy current problems based on the method of snapshots,” IEEE Transactions on Magnetics, 49(5): 1697–1700, 2013.
57. Wang, S.; Wang, H.; Bian, Z.; Chen, S.; Song, P.; Su, B.; Gao, W. A Robust Direction-of-Arrival (DOA) Estimator for Weak Targets Based on a Dimension-Reduced Matrix Filter with Deep Nulling and Multiple-Measurement-Vector Orthogonal Matching Pursuit. *Remote Sens.* **2025**, 17, 477. DOI: 10.3390/rs17030477
58. N. C. Nguyen, “Generative Reduced Basis Method,” arXiv:2410.05139 [math.NA], 2024.
59. Berkooz, G., Holmes, P., & Lumley, J. L. (1993). The proper orthogonal decomposition in the analysis of turbulent flows. Annual Review of Fluid

Mechanics, 25, 539–575. (Demonstrates that POD modes are ordered by energy content and that most flow energy is captured by the first few modes.)

60. Epps, B. P., & Techet, A. H. (2010). An error threshold criterion for singular value decomposition modes extracted from PIV data. Experiments in Fluids, 48(2), 355–367. (Shows that retaining modes beyond a certain cutoff can introduce noise, and provides a criterion to truncate modes to avoid overfitting experimental noise.)
61. ANSYS Inc. (2009). ANSYS Fluent 12.0 Theory Guide. Canonsburg, PA: ANSYS Inc.
62. Kang, Q., Feng, X., Yang, C., & Wang, J. (2022). DEM–VOF simulations on the drawdown mechanisms of floating particles at free surface in turbulent stirred tanks. Chemical Engineering Journal, 431, 133275.
63. Ting, S., Hu, Y., Wang, W., Jin, Y., & Cheng, Y. (2013). Simulation of solid suspension in a stirred tank using CFD-DEM coupled approach. Chinese Journal of Chemical Engineering, 21(10), 1069–1081.
64. De Lamotte, A., Delafosse, A., Calvo, S., & Toye, D. (2018). Analysis of PIV measurements using modal decomposition techniques (POD and DMD) to study flow structures and their dynamics within a stirred-tank reactor. Chemical Engineering Science, 178, 348–366.
65. OpenCFD Ltd. (2020). OpenFOAM v8 User Guide: Standard k–e Model Coefficients. OpenFOAM Foundation. (Retrieved from OpenFOAM Documentation).
66. Musik, M., & Talaga, J. (2016). Investigation of hydrodynamics in an unbaffled stirred vessel with an eccentrically located Rushton turbine. *Czasopismo Techniczne*.

67. Wu, M., Jurtz, N., Walle, A., & Kraume, M. (2022). Evaluation and application of efficient CFD-based methods for the multi-objective optimisation of stirred tanks. *Chemical Engineering Science*, *263*, 118109.