Upstream
Applied Science

Development and Application of an Open Access Bioreactor Performance Modelling Workflow for Cultivated Meat Yield Prediction and Optimisation

Simon Hubbard

Upstream Applied Science Ltd.

Abstract

The ability to predict the yield of cultivated meat production facilities is valuable in providing cost of goods sold predictions as part of techno-economic analyses and as a key performance metric in bio-manufacturing facility design and optimisation activities ahead of capital investment and facility construction. Existing cultivated meat techno-economic analyses either assume an achievable yield or predict it using model equations which are quick to solve but omit or simplify key performance characteristics. There is, at present, no openly available modelling workflow that enables yield prediction or optimisation for given bioreactor architecture, geometry and operational inputs. The present work builds on existing yield prediction approaches through the use of computational fluid dynamics to provide more targeted performance characterisation, coupled to a Bayesian optimisation workflow to optimise yield against bioreactor geometric and operational input parameters and made openly available. Application to 20,000 litre and 42,000 litre stirred tank reactors suggest bioreactor performance limited yield (not cell metabolism limited yield) could be higher than reported in previous studies. Application to a previously considered 300,000 litre air lift reactor design suggests that yields used in a recent techno-economic analysis are achievable, from a bioreactor performance perspective.

1 Introduction

Conventional meat production through animal agriculture has well documented sustainability and capacity challenges. Global greenhouse gas (GHG) emissions from meat production contribute to 54% of all agriculture-based emissions (OECD [2021]), which make up 17% of the global emissions (FAO [2021]), putting meat production at more than 9% of global emissions. Beef has the highest footprint, with one kilogram of final product generating 99.5 kg of CO_2 equivalent green house gas omissions and requiring 25 kg of dry animal feed. 46% of the habitable land on Earth is used for animal agriculture, of which 76% is used to provide only 37% of the global protein supply (Ritchie and Roser [2021]). Providing food for a global population predicted to grow from 8.0 billion in November 2022 to 9.7 billion in 2050 and almost 10.4 billion in the mid 2080's (United Nations and Social Affairs [2022]) will require considerable conversion of existing forest to animal agriculture. Alternative means of providing protein, either through plants or cultivating meat cells separate to animals, offer the possibility of providing protein nutrition with a significantly reduced environmental footprint and within the resources available to meet future demand. Cultivating meat, which involves growing cells in media in bioreactors of progressively larger scales, is one of the alternative means of providing protein. One of the primary challenges of cultivating meat is developing the capability to produce consumer products that compete effectively on cost with conventionally farmed meat.

Constructing models of production yield and production facility capital and operating costs is therefore extremely valuable, both for the ability to predict costs ahead of facility development and to ascertain which aspects of the relevant cell biology and industrial production may be the greatest cost limiting constraints. Such modelling presents both breadth and depth challenges, requiring attention to the details of cell metabolism and growth, fluid and chemical dynamics of bioreactor performance and bio-production facility design, construction and operation. A common approach to predicting the cost of goods sold (COGS) for cultivated meat is techno-economic analysis (TEA) which is a high level system modelling activity utilising process design and production facility performance and cost modelling. Such activities often use laboratory and/or pilot scale performance data in combination with predicted improvements with scale. For cultivated meat analyses a key output is the production yield from the bio-process, which is then combined with the capital and operating costs of the production facility to generate a cost of goods sold.

Previous TEA's for cultivated meat (Negulescu et al. [2023]; Humbird [2020]; Risner et al. [2020]; Vergeer et al. [2021]; Garrison et al. [2022]) have used a variety of tools and modelling approaches, and have significant differences of values for model inputs, the cost of media being the most apparent, and wide variation in the predicted COGS. Tables 1 and 2

Study	Line	Diameter	Hydrated Mass	Dry Mass Fraction	Doubling Time	Oxygen Uptake Rate	Respiratory Quotient	Ammonia Production Rate	Lactate Production Rate
		[μm]	[pg]		[hr]	$\left[\frac{\text{mmol}}{g_{\text{dry}}\text{hr}}\right]$		$\left[\frac{\text{mmol}}{g_{\text{dry}}\text{hr}}\right]$	$\left[\frac{\text{mmol}}{g_{\text{dry}}\text{hr}}\right]$
Vergeer et al. [2021]	Unspecified MSC	9.4	3084	0.23-0.3	21-35	-	-	-	-
Risner et al. [2020]	Unspecified MSC	-	-	-	8-24	Various	-	-	-
Humbird [2020, 2021]	CHO	17.7	3000	0.3	24	0.45 - 0.50	1.20 - 1.27	0.005 - 0.092	0.053 - 0.624
Negulescu et al. [2023]	Unspecified MSC	17.7	3000	0.3	23	Referenced	1.20	-	-

Table 1: Cell lines and characteristics considered in existing TEA models. MSC: Mesenchymal Stem Cell; CHO: Chinese Hamster Ovary.

summarise the cell line and bioreactor characteristics, respectively, used in previous TEAs. From a bioreactor performance characterisation perspective, the maximum achievable cell density is key due to it's direct link to the COGS. Some studies use a fixed value as input, others evaluate it as part of the modelling process. The TEA of Humbird, first in an open access publication (Humbird [2020]) and subsequently summarised in a journal publication (Humbird [2021]) provides the most detailed approach to determine the maximum achievable yield by modelling the mass transfer, mixing time and induced hydrodynamic stress characteristics of stirred tank bioreactors from 1,000 to 250,000 litre volumes. The open access publication is more thorough and referenced subsequently here. An overview of specific bioreactor performance models employed in previous TEAs is shown in Table 3, highlighting the extent of model use by Humbird [2020]. A more detailed review is provided in the sections below.

Study	**Line**	**Type**	**Volume (Working)**	**Doublings**	**Starting Cell Density**	**Final Cell Density**	**Oxygen Uptake Rate**	**Oxygen Transfer Rate**	**Yield**
			$[\mathrm{m}^3]$		[cells/ml]	[cells/ml]	$\left[\frac{\mathrm{mmol}}{\mathrm{L\,hr}}\right]$	$\left[\frac{\mathrm{mmol}}{\mathrm{L\,hr}}\right]$	$\mathrm{g_{wet}/L}$
Vergeer et al. [2021]	Unspecified MSC	STR & PR	0.06 (0.05), 12.5 (10), 2.5 (2)	$7-11$	-	5×10^7	-	-	
Risner et al. [2020]	Unspecified MSC	STR	20	3-10	-	1×10^7 – 2×10^8	Various	Various	
Humbird [2021, 2020]	CHO	STR	1-200	2 or 6	-	Constrained	Time varying, reaction dependent	Modelled	46, 120 (fed-batch, sparge dependent), 195 (perfusion)
Negulescu et al. [2023]	Unspecified MSC	STR & ALR	42 & 211 (STR), 262 (ALR)	2.25	6.6×10^6	3.3×10^7	≤ 10.3	Modelled	100

Table 2: Bioreactor characteristics and performance considered in existing TEA models. STR: Stirred Tank Reactor; PR: Perfusion Reactor; ALR: Air Lift Reactor.

TEA	O_2 Transport	CO_2 Transport	Mixing Time	Hydrodynamic Stress
Risner et al. [2020]	Y	N	N	N
Humbird [2020]	Y	Y	Y	Y
Vergeer et al. [2021]	N	N	N	N
Garrison et al. [2022]	N	N	N	N
Negulescu et al. [2023]	Y	N	N	N

Table 3: Bioreactor performance characteristics modelled in previous techno-economic analyses.

1.1 Oxygen Mass Transfer

For cells to grow, the overall rate at which oxygen is transferred in to the media (the oxygen transfer rate, OTR) must be equal to or greater than the rate at which it is consumed by the cells (the oxygen uptake rate, OUR). The oxygen transfer rate (OTR) is often represented as the product of a bioreactor system level mass transfer coefficient, k_La, and a gas concentration difference:

$$\mathrm{OTR} = k_La(C^*_{O_2} - C_{O_2}) \tag{1}$$

where C_{O_2} is the dissolved oxygen concentration the media and $C^*_{O_2}$ is the saturated concentration at the gas liquid interface, which can be calculated using Henry's law and the gas concentration. The concentrations are volume averages, with no consideration for spatial variation. The system level mass transfer coefficient for stirred tank bioreactors is often modelled using an equation of the form

$$k_La = A\left(\frac{P}{V}\right)^{\alpha}(u_s)^{\beta} \tag{2}$$

where P is the impeller power, V the media volume, u_s the sparged gas superficial velocity and A, α, β are model constants. Humbird [2020] uses values from a prior study (Xing et al. [2009]) of CHO cell growth in a 5,000 litre stirred tank reactor (STR) giving $A = 0.075, \alpha = 0.47, \beta = 0.8$, whereas Negulescu et al. [2023] use $A = 0.002, \alpha = 0.7, \beta = 0.2$ for STRs and $A = 0.32, \alpha = 0, \beta = 0.7$ for airlift reactors (ALRs). The difference in model constant values is apparent and likely driven by different bubble size distribution dynamics, such as coalescing and non-coalescing behaviour, in the prior works which evaluated the constants.

For the concentration difference $(C^*_{O_2} - C_{O_2})$ in equation 1, Humbird [2020] employs a log mean difference:

$$C^*_{O_2} - C_{O_2} = \frac{\left(C^*_{O_2} - C_{O_2}\right)_{\text{btm}} - \left(C^*_{O_2} - C_{O_2}\right)_{\text{top}}}{\ln\left[\frac{\left(C^*_{O_2} - C_{O_2}\right)_{\text{btm}}}{\left(C^*_{O_2} - C_{O_2}\right)_{\text{top}}}\right]} \tag{3}$$

with terms at the top (top) and bottom (btm) of the media calculated from operational characteristics and steady state mass flux balances. Negulescu et al. [2023] assume a maximum possible dissolved oxygen concentration in place of the saturated concentration (6.7 mg/L) and a set point of 2.5mg/L for the dissolved concentration.

Risner et al. [2020] use a slightly different form for the oxygen transfer rate:

$$\text{OTR} = kA(C^*_{O_2} - C_{O_2}) \tag{4}$$

where the volumetric mass transfer coefficient $k_L a$ is replaced with the coefficient k and the mean bubble area A, varied as part of a broad sensitivity analysis.

1.2 Carbon Dioxide Mass Transfer

In addition to transferring oxygen from the gas to the media, the bioreactor must also remove the carbon dioxide produced by the cells from the media to the gas. An analogous approach to that used for oxygen transfer could be applied to evaluate the carbon dioxide removal rate, however model equations for the carbon dioxide mass transfer coefficient are not as prevalent as for oxygen, and published measurements of dissolved carbon dioxide concentrations in production scale bioreactors are rare. Given this lack of relevant data and models, Humbird [2020] considers the saturated concentration of dissolved carbon dioxide at the top of the media based on the mole fraction of carbon dioxide being transferred in to the headspace from a mass balance between the injected sparged gas and the gas consumption and production by cells. The dissolved carbon dioxide concentration in the media that is in interfacial equilibrium with the headspace concentration is

$$pCO_2 = \gamma_{CO_2} p_b \tag{5}$$

where γ_{CO_2} is the mole fraction of carbon dioxide in the headspace and p_b is the headspace back pressure. This dissolved carbon dioxide concentration is then required to be below the limit of 100 milli-bar based on prior studies of concentrations that limit growth in CHO cells (Gray et al. [1996]).

1.3 Mixing

In addition to providing sufficient oxygen to the cells, the bioreactor must also ensure that it is well distributed throughout the media (on the assumption that the cells are also well distributed). Humbird [2020] addresses this requirement by utilising a model equation for the mixing time of stirred tank reactors from Nienow [2014]:

$$t_m \approx 6T^{2/3}\left(\frac{P}{V\rho_L}\right)^{-1/3}\left(\frac{D}{T}\right)^{-1/3}\left(\frac{H_L}{T}\right)^{2.5} \tag{6}$$

where ρ_L, D, H_L and T are the media density, impeller diameter, liquid height and bioreactor diameter respectively. To ensure well distributed dissolved oxygen Humbird [2020] requires that the mixing time be less than the mass transfer time:

$$t_m < \frac{1}{k_L a} \tag{7}$$

1.4 Induced Hydrodynamic Stress

The hydrodynamic environment inside a bioreactor has an impact on cell behaviour and may cause the cells to reduce the rate at which they grow or to cause them to die. Existing characterisations of this imposed hydrodynamic stress include shear stress, shear rate and the rate at which kinetic energy is dissipated to heat at the smallest scales of turbulent motion. The energy dissipation rate has been characterised both as the average and maximum rates within the media, with the maximum dissipation rate occurring in close proximity to the impeller (Chalmers [2015, 2021]; Nienow [2006]).

$$\lambda_k = \left(\frac{\nu^3}{\varepsilon_{\max}}\right)^{\frac{1}{4}} \tag{8}$$

Humbird [2020] follows a commonly used approach of converting the peak dissipation rate into a Kolmogorov length scale, equation 8, which is then compared to the cell diameter to determine if the cell experiences any of the dissipated energy; it is proposed that if the cell diameter is less than the Kolmogorov length scale then it does not experience any energy dissipation as it is occurring at scales larger than the cell, and it's growth is unaffected by the hydrodynamics. This approach has both theoretical and practical drawbacks. From a theoretical perspective, Kolmogorov's hypothesis considers that the flow structure is determined by the average energy dissipation rate, not the peak, and the length scale at which the energy dissipation peaks is an order of magnitude greater than the Kolmogorov length scale (Pope [2000]). From a practical perspective, the peak energy dissipation rate used in the calculation is taken as a time averaged value, whereas the cells moving through the media around the impeller will experience the instantaneous local energy dissipation rates, not a time averaged value. Such drawbacks may be responsible for poor correlation between the Kolmogorov length scale, the cell diameter and cell growth limitation or death (Chalmers [2021]).

Despite such challenges, consideration of the effects of hydrodynamic stress on cell behaviour is critical to predict yield. Humbird [2020] uses a minimum viable Kolmogorov length scale of $20\mu m$, equivalent to the considered CHO cell diameter, to determine the limiting peak energy dissipation rate, and to subsequently determine the limiting average energy dissipation rate through a relationship between the maximum, $\varepsilon_{\max}$ and average, $\overline{\varepsilon}$, energy dissipation rates for Rushton turbines, $\varepsilon_{\max} = 48\overline{\varepsilon}$, from Nienow [2009]. Such calculation results in limiting the impeller power so that the average power per unit volume (P/V in equation 2) does not exceed 64 W/m^3. By comparison, for 20,000 litre and 42,000 litre STRs, Negulescu et al. [2023] use $P/V = 500$ W/m^3, an order of magnitude higher. The limit in Humbird [2020] is likely to be overly restrictive as a result of applying the relationship between average and maximum dissipated power from a single impeller to STRs with two impellers. When the volume averaged power is taken as being provided by two impellers, the limit becomes 128 W/m^3.

The motion of bubbles through the media and bursting at the free surface also generates energy which is dissipated to heat, and can be assessed in a similar way. Considering the energy dissipated from the relative motion of bubbles through the media, the conditions can also be characterised in terms of an eddy dissipation. Such evaluation is complex and experimental experience offers a simper consideration: Humbird [2020] uses a limit of a sparged gas superficial velocity at the top of the vessel of $u_s = 0.006$ m/s, equivalent to a gas sparging rate of 0.1 vessel volumes per minute (vvm) in a 20,000 litre STR. Negulescu et al. [2023] use a limit of 0.01 vvm for all bioreactor volumes considered.

1.5 Opportunities for Bioreactor Performance Modelling

Prior TEAs have used assumed values for a range of performance aspects, with Humbird's work providing the most detailed technical modelling framework. This framework enabled a quick optimisation of bioreactor operating conditions to be performed in Excel, providing predicted optimum maximum yields for STRs of a wide range of working volumes. A drawback to this approach is the description of bioreactor performance with empirical equations from prior works at specific working volumes. The constants used in Equation 2 are from work using a 5,000 litre STR, which are then applied across a much wider range, 1,000 to 250,000 litres, with no consideration of other critical aspects of performance, such as the size of the bubbles generated by the sparger.

It would be valuable to use a similar workflow to enable optimisation of both bioreactor operation and geometry, allowing consideration of bespoke impeller geometries and bubble size distributions, which may provide scope for generating higher maximum yields than expected from the more constrained approaches used so far. In this paper, workflows are developed that enable such optimisation by using computational fluid dynamics (CFD) to provide performance characteristics for bioreactor operational and geometric variation. Two bioreactor "system level" models are used, one following the approach of Humbird [2020] but with the use of CFD generated performance metrics, and the second a new approach that aims to improve on some aspects of the performance modelling, while still using CFD generated performance metrics. Figure 1 illustrates these workflows schematically.

In addition, CFD models have the ability to predict the levels and effects of spatial variability of dissolved oxygen and nutrients on cell growth and hence maximum yield. A CFD modelling approach is used here to provide insight in to the likelihood of spatial variability effecting the predicted maximum yield from the system level models, which assume spatial homogeneity in order to enable simpler and more timely yield calculation.

These system level and CFD models are applied to the optimisation of 20,000 and 42,000 litre STRs in fed batch model as considered by Humbird [2020] and Negulescu et al. [2023] respectively. The CFD model for assessing spatial variability is applied to a 300,000 ALR as assessed previously in CFD by Li et al. [2020], and close to the volume of the ALR considered by Negulescu et al. [2023].

The focus of these models is to provide insight in to the yield potential achievable through bioreactor design and operation. To this end, yield limits associated with metabolite concentrations are not considered in the models, as these are indepen-

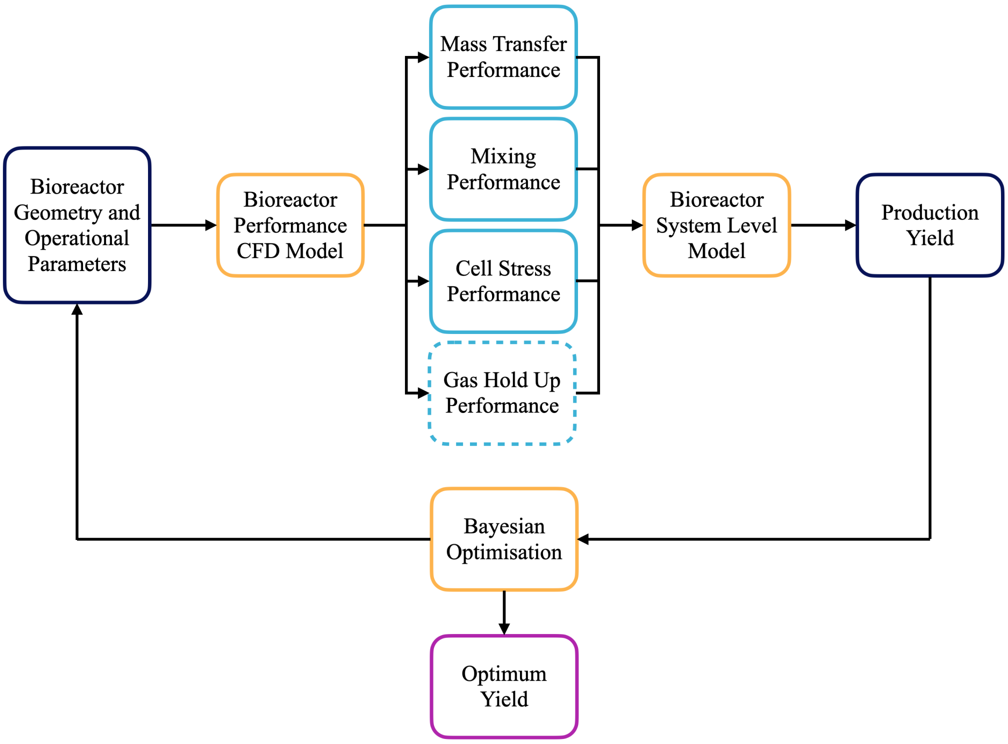

Figure 1: Optimisation workflows developed in the present study. Bioreactor inputs (geometry, operating parameters) and outputs (yield) in dark blue, models (CFD, bioreactor system and optimisation) in orange, bioreactor performance characteristics in cyan and predicted optimum yield in purple. The additional gas hold up performance used in system model B is highlighted with a dashed box.

dent of bioreactor mixing, mass transfer and induced stress performance. The bioreactor yield potential considered here could subsequently be combined with the yield potential from improved cell response to metabolite concentrations that may be achieved through cell line selection or metabolic engineering.

2 Methods & Materials

2.1 Bioreactor System Level Models

System level bioreactor models neglect spatial variability within the media, which allows them to predict performance over the whole production duration, which is not possible with full three dimensional CFD models that capture spatial variability due to the computational effort required. Two system level models are used, the first based on the approach of Humbird [2020], the second a new approach which aims to improve the representation of carbon dioxide mass transfer by making use of an additional bioreactor performance characteristic, gas hold up.

2.1.1 System Level Model A

The yield is evaluated from the number of cells in the media at the time when cell growth is limited by one of the constraints. The number of cells over time is evaluated from:

$$N(t) = N_0 \, e^{(\mu t)} \tag{9}$$

where N_0 is the initial number of cells, μ the growth rate and t time. The mixing time, mass transfer, carbon dioxide concentration and induced stress constraints are as follows.

The mixing constraint, where the mixing time must be less than the mass transfer time is implemented as in equation 7, with the mixing time, τ_m and oxygen mass transfer coefficient, $k_L a$, evaluated from a CFD model. This constraint is independent of the process duration, t, in equation 9.

The oxygen mass transfer constraint is implemented by limiting the process duration time, t, in equation 9, to times when the total oxygen uptake rate of all cells is less than or equal to the rate at which oxygen is being transferred in to the media.

This can be expressed as:

$$k_L a(C^*_{O_2} - C_{O_2}) \geq q_{O_2} N(t) \tag{10}$$

The mass transfer coefficient, $k_L a$, is evaluated from a CFD model, the concentration difference, $(C^*_{O_2} - C_{O_2})$, is evaluated from equation 3, and q_{O_2} is the oxygen uptake rate per cell.

The carbon dioxide concentration constraint is implemented by limiting the process duration time to conditions when the dissolved carbon dioxide concentration is less than or equal to a prescribed limit:

$$C_{CO_2} \leq C^{\lim}_{CO_2} \tag{11}$$

where C_{CO_2} is evaluated from equation 5 and $C^{\lim}_{CO_2}$ is taken to be 100 millibar as in Humbird [2020]. The mole fraction of carbon dioxide in the headspace as required in equation 5 is evaluated as in Humbird [2020].

The impeller induced stress constraint is implemented by limiting the maximum time averaged turbulence energy dissipation rate to be less than or equal to a prescribed limit:

$$\overline{\varepsilon_M} \leq \overline{\varepsilon_M}^{\lim} \tag{12}$$

The time maximum time averaged turbulence energy dissipation rate is evaluated from a CFD model, and the limiting value is prescribed as an input. The bubble motion and rupture stress limit is implemented as:

$$u_s \leq u_s^{\lim} \tag{13}$$

The superficial velocity, u_s, and it's limit are evaluated as in Humbird [2020], with $u_s^{\lim} = 0.006$ m/s.

The oxygen uptake rate per cell is taken to be 0.490 millimoles per hour per gram dry weight and the carbon dioxide production rate is 0.593 millimoles per hour per gram dry weight, corresponding to the values derived by Humbird [2020] for the intermediate CHO metabolism. The intermediate metabolism is chosen here as a compromise between the wild-type metabolism and the enhanced metabolism which was developed to take advantage of sparging gas with 90% oxygen. To be consistent with this intermediate metabolism, all CFD models are run with 21% oxygen gas. Figure 2 shows the results from system level model A when run with the wild-type and intermediate metabolisms, showing visually identical results to Figure 2.7 from Humbird [2020] for the two doublings scenario.

2.1.2 System Level Model B

System level model A is limited by it's approach to carbon dioxide; it does not consider CO_2 mass transfer as rigorously as it does for oxygen, and the concentration constraint is based on an equilibrium saturated concentration in the liquid at the headspace free surface rather than a concentration in the whole liquid media volume. System level model B improves these aspects by utilising a carbon dioxide mass transfer coefficient and the steady state bubble volume both evaluated from a CFD model.

In model B, the volume average gas phase (bubbles) and liquid phase (media) concentrations of both oxygen and carbon dioxide are described by coupled ordinary differential equations:

$$\frac{dC_l}{dt} = k_L a(C_g H - C_l) + qN \tag{14}$$

$$\frac{dC_g}{dt} = -k_L a(C_g H - C_l) + Q(C_g^{in} - C_g)/V \tag{15}$$

The equation for dissolved concentration in the liquid media (14) considers the transfer from the gas through the $k_L a(C_g H - C_l)$ term, where H is Henry's coefficient for the species concerned, and the consumption by the cells through the qN term, where q is the cell uptake or production rate and N the number of cells. The equation for the concentration in the gas bubbles (15) considers the transfer to the liquid through the $-k_L a(C_g H - C_l)$ term, and the increase in concentration due to freshly sparged gas through the $Q(C_g^{in} - C_g)/V$ term, where Q is the volumetric sparging rate, C_g^{in} the concentration in the sparged inlet, and V the steady state volume of immersed gas bubbles.

Equations (14, 15) are solved using SciPy's odeint function, taking mass transfer coefficients for oxygen and carbon dioxide and a steady state immersed bubble gas volume from a CFD model, with all other parameters specified as direct model inputs.

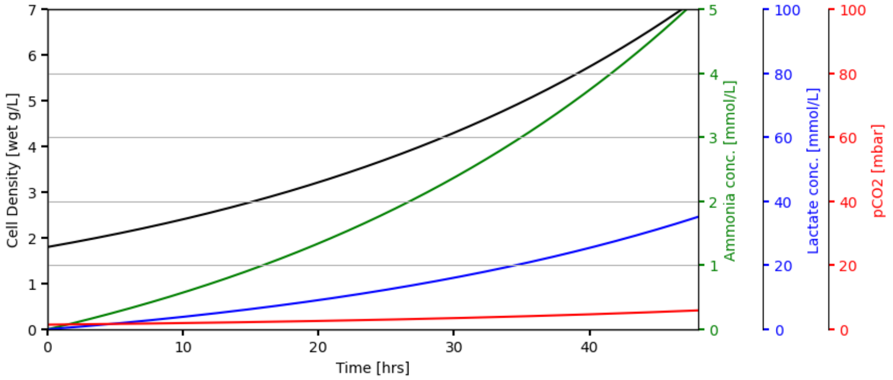

8

(a) Lac/Glc=1.77, Gln/Glc=0.13

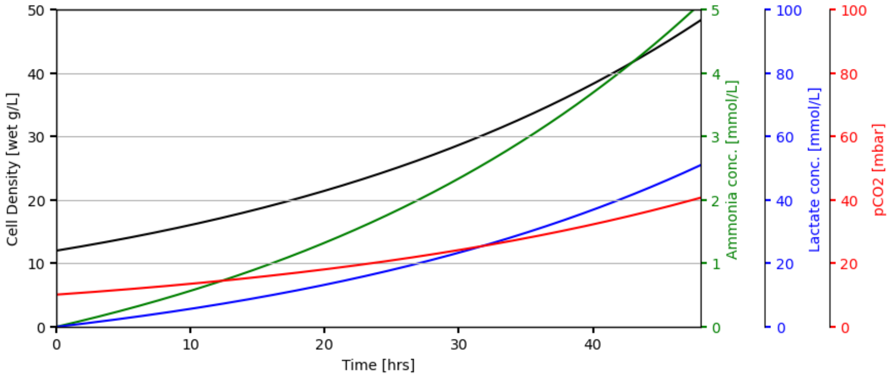

(b) Lac/Glc=0.94, Gln/Glc=0.047

Figure 2: Time evolution of cell density and metabolite concentrations from the present study for the wild-type (a) and intermediate (b) CHO cell metabolisms from Humbird [2020].

The mixing constraint now considers that the mixing time must be shorter than both the oxygen and carbon dioxide mass transfer times:

$$\tau_m < \frac{1}{k_L a^{O_2}} \quad \text{and} \quad \tau_m < \frac{1}{k_L a^{CO_2}} \tag{16}$$

The mass transfer constraint now limits the production duration to times at which dissolved oxygen is present and dissolved carbon dioxide is lower than the prescribed limit:

$$C_l^{O_2} > 0 \quad \text{and} \quad C_l^{CO_2} < \left(C_l^{CO_2}\right)^{\text{lim}} \tag{17}$$

with the CO_2 limit at 100 millibar as used in model A.

The impeller induced stress constraint and oxygen uptake and carbon dioxide production rates are as used in model A. Model B does not implement a bubble motion hydrodynamic stress constraint through limiting the superficial gas velocity as in model A. The oxygen uptake and carbon dioxide production rates are as used in system level model A.

2.2 Lattice Boltzmann Computational Fluid Dynamics

Computational fluid dynamics simulations of bioreactors face a number of challenges, across turbulence modelling and transport, bubble transport and mass transfer, and computational efficiency. Lattice Boltzmann approaches are well suited to graphics processing unit (GPU) architectures, enabling similar computational efficiency on a small number of GPUs to that found for finite volume approaches on high performance compute clusters comprised multiple many core central processing units. The computational efficiency and physical accuracy benefits of the commercial lattice Boltzmann CFD package M-Star in simulating bioreactors has been shown by Haringa [2022]. M-Star is used here to provide two simulations: 1) the prediction of bioreactor performance characteristics such as mass transfer, hydrodynamics stress, gas holdup and mixing time, and 2) the prediction of the effects of spatial variation in dissolved oxygen concentration on the achievable yield through coupling with a simple metabolic model of dissolved oxygen concentration dependent oxygen uptake rate. M-Star is run on two NVLinked NVidia RTX A6000 GPU's, providing 80 teraflops of single precision compute capability and 96GB RAM.

M-Star uses the D3Q19 velocity vector set for discretisation of the velocities (M-Star). The equations are solved on a fixed computational lattice, analogous to cubed finite volumes, with the spatial resolution defined as the number of lattice elements used to discretise the x-direction extent of the bioreactor, equal to the diameter. Stationary boundaries are grid aligned, with test simulations replicating previous results (Haringa [2022]) showing negligible effect on the outputs of varying between grid aligned and wall interpolated boundaries. Moving bodies are represented using the immersed boundary method (Zhu et al. [2011]). All walls are no-slip surfaces, with the top domain boundary taken to represent the free surface between the liquid and gas and represented by a no-shear condition. Such an implementation does not resolve the free surface, which enables use of larger timesteps to reduce the computational effort. The consequence of this approach is that the mass transfer from the headspace gas in to the media is not captured. This interface mass transfer is negligible compared to the mass transfer from sparged bubbles at sufficiently large bioreactor volumes ($\geq 5,000$ litres).

Turbulence is modelled using a Large Eddy Simulation approach, with the Smagorinsky subgrid scale model. All simulations are performed with the same value of the Smagorinksy subgrid model constant, $C_{LES} = 0.1$. Wall functions are not considered, due to their effect being shown to be small in the context of comparison to experimental impeller discharge profiles (Haringa [2022]).

In addition to solving for the turbulent motion of the liquid phase, the transport of an injected dye and dissolved gases within the liquid media are modelled using the general scalar transport equation

$$\frac{\partial C_s}{\partial t} = \nabla \cdot (D_s \nabla C_s) - \nabla \cdot (u C_s) \tag{18}$$

where C_s is the concentration, D_s the mass diffusivity of the species concerned and u is the velocity of the media.

The motion of gas bubbles injected from a sparger is captured by a Lagrangian approach. The bubbles are removed at the no-shear boundary at the top of the media to represent the transmission to the headspace gas. In order to reduce the computational effort required to model a large number of bubbles, a parcel approach is utilised in which the Lagrangian dynamics of parcels, each representing a number of bubbles, is computed using Newton's second law (Thomas et al. [2022a]):

$$m_p a_p = F_{g,p} + F_{a,p} + F_{d,p} \tag{19}$$

where m_p, a_p are the parcel mass and acceleration, and $F_{g,p}$, $F_{a,p}$, $F_{d,p}$ are the forces on the parcel from gravity, added virtual mass and drag respectively. The parcels are two way coupled to the fluid using Newton's third law, based on either just the buoyancy force $g\alpha(\rho_g/\rho_f - 1)$ or all of the forces in equation 19. The use of the simpler buoyancy force coupling offers a reduction in computational effort through the use of larger time steps.

The mass transfer of gases between the parcels and the media is modelled by

$$\dot{n} = k_L^b A_p (HC_g - C) \tag{20}$$

where $\dot{n}$ is the mass flux from the parcel, k_L^b is the bubble specific mass transfer coefficient and A_p is the parcel surface area:

$$A_p = N_p \pi d_p^2 \tag{21}$$

where N_p is the number of bubbles in the parcel and d_p is the bubble diameter, which changes in response to the local hydro-static pressure through

$$d_b = d_{b,P_H} \left(\frac{P_H}{P_H + \rho_L g h_b} \right)^{(1/3)} \tag{22}$$

where d_{b,P_H} is the diameter at the headspace pressure P_H, ρ_L the media density, g gravitational acceleration and h_b the depth of the bubble below the head space free surface. The effects of coalescence and break up of bubbles and bubble parcels can be captured through consideration of a critical coalescence Reynolds number and the ratio of local energy dissipation rate to bubble surface energy (Thomas et al. [2022b]). Coalescence and splitting is ignored in the present study, in alignment with previous studies where the presence of surfactant and anti-foaming agents suppresses such behaviour (Li et al. [2020]).

The bubble specific mass transfer coefficient, k_L^b, can be represented by an empirical equation, or one derived from a theoretical analysis. Previous studies with M-Star shows that using a constant value can provide accurate predictions with reduced computational cost (Farsani et al. [2022]). The approach of Kawase and Moo-Young [1990] is used here:

$$k_L^b = 0.3(D_{L,O_2})^{\frac{1}{2}} \left(\frac{\varepsilon}{\nu}\right)^{\frac{1}{4}} \tag{23}$$

where D_{L,O_2} is the mass diffusivity of oxygen in the media and ε the local eddy dissipation rate.

2.2.1 Validation

Previous work has validated M-Star for bioreactor mass transfer (Thomas et al. [2021b, 2022b, 2023]; Farsani et al. [2022]; Schlaich et al. [2023]), mixing time (Thomas et al. [2021a]; Fitschen et al. [2021]) and impeller power number (Thomas et al. [2021a]; Haringa [2022]; Kuschel et al. [2021]). With prior validation covering CFD predictions of mass transfer and mixing time, we are interested in further validation of the prediction of the flow field in the proximity of the impeller, where the peak energy dissipation rate occurs which drives the induced stress metric, and the gas hold up which is used in system level model B to provide the steady state immersed bubble volume.

A recent comparison of M-Star to experimental velocity and turbulence data in the discharge region of a Rushton impeller was made by Haringa [2022], and is extended here to consider higher spatial resolutions. The geometry of a bioreactor with a single Rushton turbine as described by Jahoda et al. [2007] is used with variations in impeller blade thickness, B, an often neglected dimension in mesh resolution and validation studies, the number of elements across the bioreactor diameter and blade thickness and the Courant Number. All models are run for 60 seconds of operation, reaching steady state within 20 seconds, and reporting values that are time averaged over the last 40 seconds. Table 4 shows the range of variations considered and the resulting predictions of power number and the minimum and maximum lattice densities, which are a predictor of error and should be as close to unity as possible. The number of lattice elements across the bioreactor diameter is varied to a higher value than considered previously.

Case	B [mm]	N_X	N_B	Co	$N_{P,\tau}$	$N_{P,\varepsilon}$	$\rho_{LB,\mathrm{min}}$	$\rho_{LB,\mathrm{max}}$	t [mins]
A1	5	200	3.4	0.05	3.67	4.25	0.993	1.007	1.17
A2	5	400	6.9	0.05	3.67	3.96	0.990	1.005	14.14
A3	5	600	10.3	0.05	3.57	3.96	0.987	1.005	67.66
B1	1.6	200	1.1	0.05	5.06	5.62	0.989	1.005	1.16
B2	1.6	400	2.2	0.05	3.36	4.00	0.991	1.004	14.12
B3	1.6	600	3.3	0.05	3.98	4.86	0.986	1.004	67.64
B4	1.6	800	4.3	0.05	4.38	4.95	0.984	1.005	179.78
C1	1.6	200	1.1	0.025	4.68	5.34	0.996	1.002	2.31
C2	1.6	400	2.2	0.025	3.44	4.09	0.997	1.001	23.21
C3	1.6	600	3.3	0.025	4.54	5.13	0.996	1.001	72.44

Table 4: The range of blade thicknesses, B, lattice divisions across the tank diameter, N_X, lattice divisions across the impeller blade thickness, N_B, and Courant numbers (Co) considered. The outputs are the turbine torque and energy dissipation rate derived power numbers, $N_{P,\tau}$ and $N_{P,\varepsilon}$ respectively, the minimum and maximum lattice densities, $\rho_{LB,\mathrm{min}}$ and $\rho_{LB,\mathrm{max}}$ and the simulation time for one second of bioreactor operation, t.

Table 4 shows a significant range of predicted impeller power numbers, from 3.44 to 5.62, which is wider than that reported by Haringa [2022]; 3.69 to 4.81. Experimental data has power numbers in the range 4.6 to 6.0 (Kaiser et al. [2017]; Distelhoff et al. [1995]). The range of minimum and maximum lattice densities is very similar, 0.984 to 1.007 in the present study, and 0.990 to 1.004 from Haringa [2022].

Comparison of the predicted time averaged radial velocity, turbulence kinetic energy and turbulence energy dissipation rate along a radial line at the centre of the impeller, analogous to Figure 1 from Haringa [2022] is shown in Figure 3. The trends and comparison between the CFD and experiments are similar to those in Haringa [2022], with the same characteristics and mechanisms at play. Of particular note in the present comparison are the $N_X = 600$ and 800 cases, which extend the resolution to significantly higher levels - the maximum considered in Haringa [2022] is $N_X = 360$. The higher resolutions considered here do not show significant improvement in the comparison to experimental data.

The comparison of predicted power number and discharge flow to experimental data and previous studies using M-Star indicates that employing a higher spatial resolution does not offer significant improvements in the accuracy of the predictions. The higher resolutions do require significantly longer simulation duration, and are therefore not considered suitable for use in optimisation workflows at this stage as a result of taking orders of magnitude longer to generate predictions that are not substantially more valuable in terms of predictive accuracy. As a result the $N_X = 200$, Co = 0.05 setup is selected for subsequent use. It may be valuable to note that the resolution sensitivity in terms of the number of elements across the blade thickness appears to suggest that a resolution of more than seven elements across the blade thickness leads to a mesh independent solution, which was not achieved with the highest spatial resolution with the 1.6mm thick blades. The bioreactor geometries considered subsequently will use common blade thicknesses, and the overall resolution required to achieve the indicated resolution independent solution would not be computationally efficient enough for the use here in an optimisation workflow. Thorough comparison of lattice Boltzmann simulations of Rushton (and other) impellers to experiment in terms of N_X and N_B may be valuable to determine how to best evaluate and characterise spatial resolution independence, and the most appropriate model configuration and hardware set-up and to achieve. Whilst potentially valuable, such study is outside the scope of the present work.

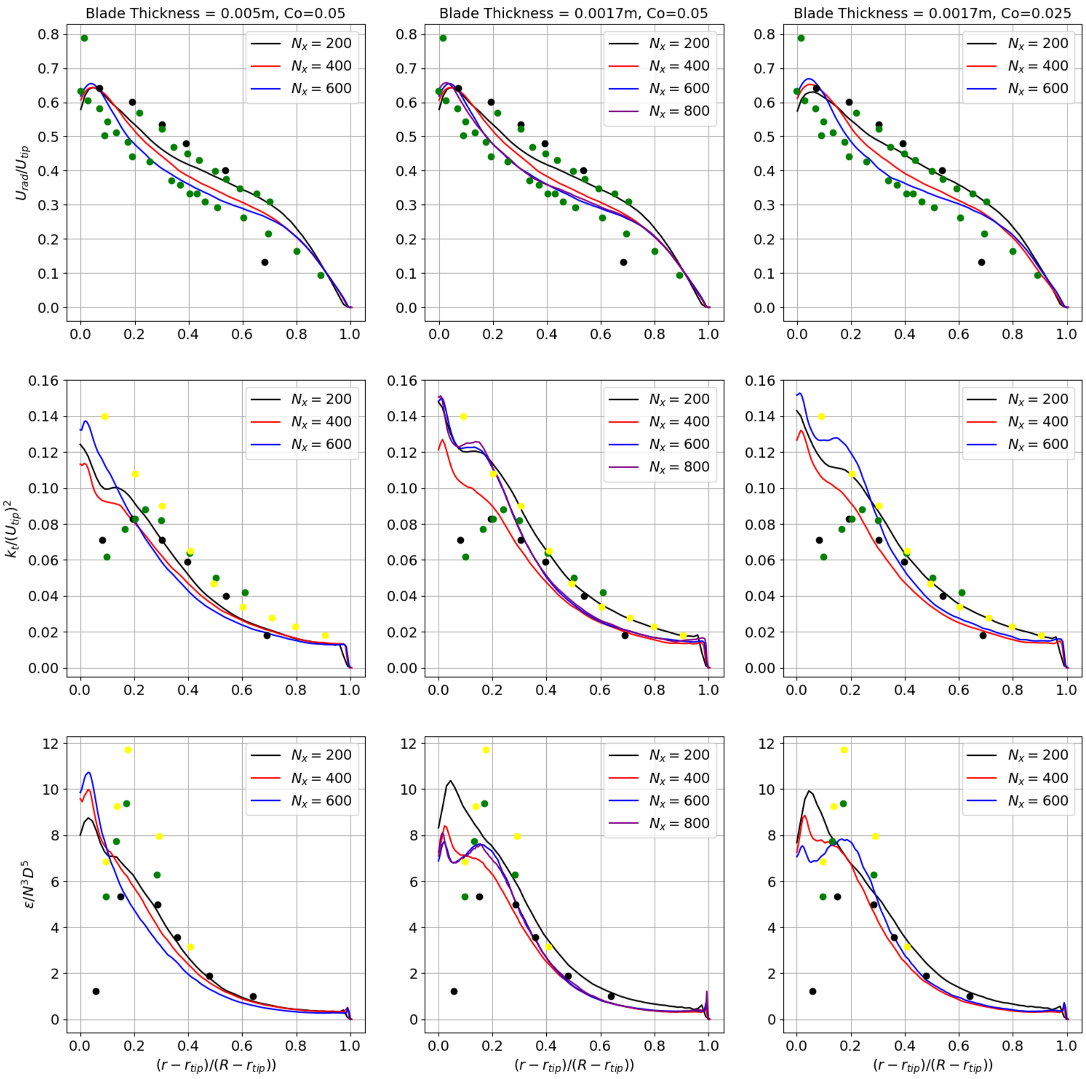

Figure 3: Resolution dependency of velocity, turbulence kinetic energy and eddy dissipation rate with blade thickness and Courant number, compared with experimental data. Top row: green circles - velocity summary by Ranade and Joshi [1990]; black circles - Haringa et al. [2018]. Middle row: green circles - Wu et al. [1989]; Wu and Patterson [1989]; ; black circles - Haringa et al. [2018]; yellow circles - Murthy and Joshi [2008]. Bottom row: green circles - Ducci and Yianneskis [2005]; black circles - Wu et al. [1989]; Wu and Patterson [1989]; yellow circles - Ducci and Yianneskis [2005].

For use with system level model B, the CFD model is required to predict the steady state immersed bubble volume, or gas hold up when expressed as a fraction of the liquid volume. As there has been less prior work to validate M-Star for gas hold up predictions, a comparison is made here with the experimental data for gas hold up in bubbles columns with a fine pore sparger as reported by Kanaris et al. [2018]. The comparison is made for an air-water bubble column of 9cm width.

Bubble splitting and coalescence is not modelled in M-Star, rather the inlet bubble size distribution is taken from the measured distribution, Figure 3.a in Kanaris et al. [2018]. The bubbles are modelled individually with no use of parcels and the bubble fluid momentum coupling is through Newton's third law with the use of equation 19. Figure 4 compares the gas hold up from Kanaris et al. [2018] with predictions from M-Star, showing good agreement in the homogeneous bubble regime, $u_s \leq 0.024$, the limit of which is an order of magnitude greater than the sparged velocity limit applied by Humbird [2020] ($u_s \leq 0.006$).

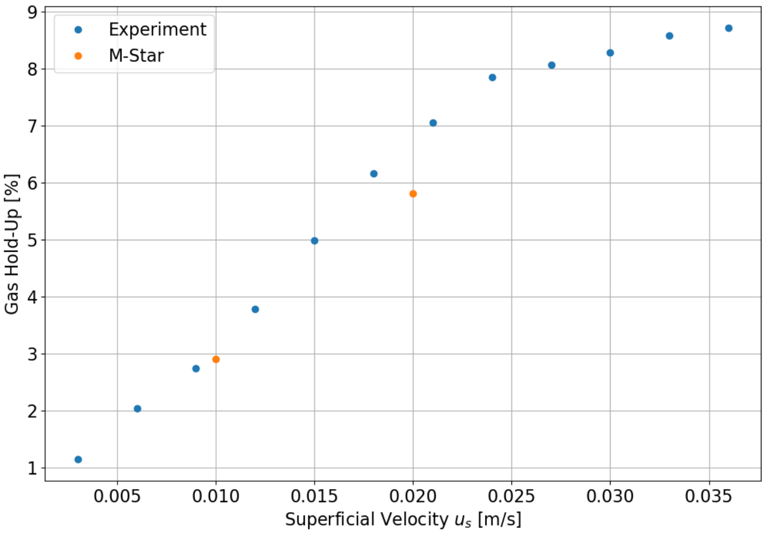

Figure 4: Comparison of predicted gas hold up from M-Star to the experimental data of Kanaris et al. [2018].

2.2.2 Bioreactor System Level Characterisation

The first CFD simulation approach provides bioreactor performance metrics - gas hold up, mixing time, mass transfer coefficients and hydrodynamic stress - which are subsequently used as inputs to system level models A and B which then subsequently predict yield. Such a combination provides a modelling workflow that takes bioreactor design and operation as inputs, and provides production yield as an output (Figure 1).

The mixing time is evaluated by depositing a conserved scalar, or dye, in to a region of the media after 10 seconds of bioreactor operation and tracking the coefficient of variation of the dye concentration over time. The time taken for the coefficient of variation to reduce to 5% is taken as the mixing time (reduction to 5% and 1% are common in the literature).

The gas hold up, or steady state immersed bubble volume is reported directly by M-Star.

The evaluation of the bioreactor mass transfer coefficients for oxygen and carbon dioxide follow the gassing out method (Garcia-Ochoa et al. [2010]). For the absorption of oxygen from the gas to the liquid the initial concentration in the media is set to zero, and the time evolution of the average concentration is used to calculate $k_L a$ from

$$\ln\left(1-\frac{\bar{C}}{C^*}\right) = -k_L a\, t \tag{24}$$

where C^* is the equilibrium interface oxygen concentration calculated from Henry's law. For the desorption of carbon dioxide from the media to the gas, the initial concentration is set to 4.464×10^{-3} moles per litre, and the time evolution of the average concentration is used to calculate $k_L a$ from

$$\ln\left(\frac{C_0}{\bar{C}}\right) = k_L a\, t \tag{25}$$

where C_0 is the initial carbon dioxide concentration. The average concentration in each equation, $\bar{C}$, is taken over values at probe locations spanning the height of the media. Probe locations for a 20,000 litre STR are shown in Figure 5. Probe location in the 42,000 litre STR model are scaled from the 20,000 litre locations.

The hydrodynamic stress is evaluated as the maximum value of the time averaged eddy dissipation rate, calculated through a custom field function in M-Star.

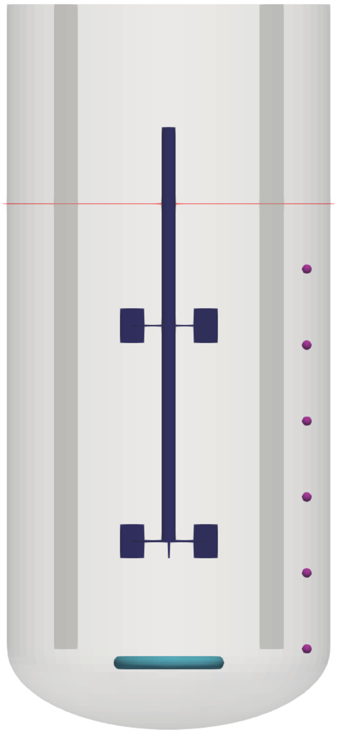

Figure 5: Probe locations (purple) in a 20,000 litre STR CFD Model. Impeller and shaft (blue), sparger (cyan) and media height (red line) are also shown.

System level characterisation CFD models couple the bubble parcels to the fluid flow using the buoyancy force only in order to reduce computational effort.

2.2.3 Bioreactor Spatial Distribution Performance

The second CFD simulation approach aims to provide an initial insight in to the effects of spatial variability on the achievable yield. The system level characterisation coupled with a CFD characterisation provides a prediction of the maximum achievable yield, assuming that the effects of spatial in-homogeneity are negligible, i.e. there are no local regions starved of oxygen. The CFD model can be set-up to predict the occurrence of such regions when coupled to cell metabolic rates. Such models cannot be run to simulate the whole production duration as they would take longer to run than the production duration. They can, however, be used to predict if local regions of low oxygen will exist at a given cell density, and therefore can be used to indicate if spatial in-homogeneity could affect a maximum achievable yield predicted by the system level characterisation and performance model workflow.

In order to capture spatial variation in dissolved gas concentration and the effects on cell growth, the general equation for the transport of a scalar used to model transport of dissolved oxygen is modified to:

$$\frac{\partial C_s}{\partial t} = \nabla \cdot (D \nabla C_s) - \nabla \cdot (u C_s) - q(C) N \tag{26}$$

where $q(C)$ is the uptake rate of the dissolved gas under consideration, which depends on the concentration, and N is the cell number density.

The concentration dependent oxygen uptake rate for CHO cells is taken from Meuwly et al. [2006]:

$$q(C_{O_2}) = q_{\max} \frac{C_{O_2}}{C_{O_2} + K_m} \tag{27}$$

where $q_{\max} = 9.44.10^{-17}$ moles per cell per second from Meuwly et al. [2006] at 37.5°C and $K_m = 2.5.10^{-3}$ moles per cubic metre from prior studies (Villadsen et al. [2011]; Miller et al. [1987]). This oxygen concentration sensitive uptake rate is compared to the constant (concentration insensitive) rates of the intermediate metabolism of Humbird [2020] and that used in the CFD analysis of a 300,000 litre ALR by Li et al. [2020] in Figure 6.

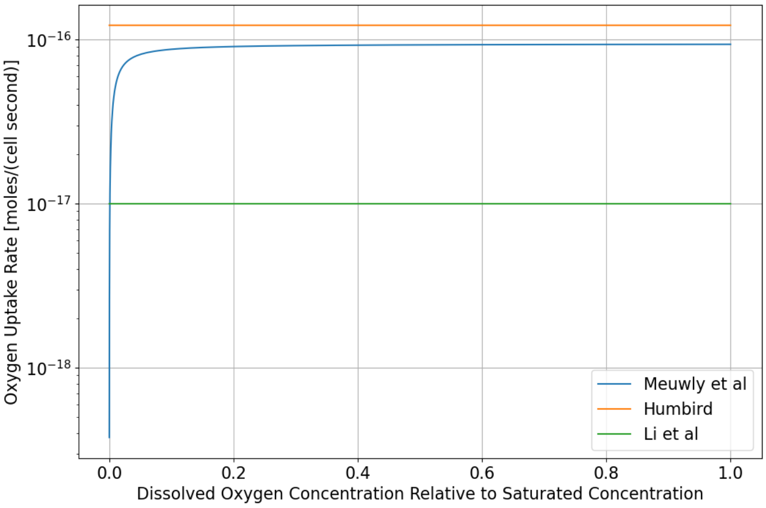

Figure 6: Comparison of oxygen uptake rates used by Meuwly et al. [2006], Humbird [2020] (intermediate metabolism) and Li et al. [2020].

Spatial distribution CFD models couple the bubble parcels to the fluid flow using the all of the forces in equation 19 as these models are single runs rather than part of an optimisation loop, where the increase in physical accuracy takes priority over the increase in model run time.

2.3 Bayesian Optimisation

The combination of a system level model and a CFD characterisation model to provide system level performance characteristics offers the possibility of optimising for the maximum achievable yield for a range of bioreactor geometries and operating parameters (see Figure 1). Bayesian optimisation is an approach that is well suited to optimising in situations where there is a significant cost, computational effort in this case, to the activity of evaluating performance. Bayesian optimisation is comprised of two key steps, the first is the use of Gaussian processes to fit a predicted performance over the input parameter space using existing data, the second is to use an acquisition function to determine the combination of parameters most likely to provide an improvement over the current optimum. Bayesian optimisation works well with CFD and other computational models as the workflow can be set-up to be self-sufficient where the CFD model is used to evaluate the performance at the conditions generated by the acquisition function, the Gaussian process response is updated and the acquisition function generates a new parameter set to evaluate, looping over a set number of iterations. Bayesian optimisation has been coupled with CFD across a range of applications, such as the optimisation of chemical (Park et al. [2018]) and fluidised bed (Kim et al. [2022]) performance.

The PyTorch based BOTorch python library provides an extensive Bayesian optimisation capability, with full exposure to all technical aspects. The Adaptive Experimentation Ax platform provides a simple python based framework to set-up and execute Bayesian optimisation workflows using the BOTorch library and is used here to automate the optimisation of bioreactor performance.

3 Results

The following optimisations and performance assessments are reported in this section.

1. Optimisation of the yield from a 20,000 litre stirred tank reactor through changes in Rushton impeller diameter, impeller blade width and RPM, and gassing rate. A number of optimisations are performed to illustrate the effects of other parameters, such as the hydrodynamic stress limit and bubble size, and the effects of the bioreactor system level model.
2. Optimisation of the yield from a 42,000 litre stirred tank reactor through changes in Rushton impeller diameter, impeller blade width and RPM, and gassing rate.
3. Assessment of the potential impact of spatial variability in oxygen availability in optimised designs from 1. and 2.

4. Assessment of the potential impact of spatial variability in oxygen availability in a previously studied 300,000L air lift bioreactor at previously considered cell densities.

Python code for the system level models, workflows, results and M-Star CFD models are available in a GitHub repository.

3.1 20,000 Stirred Tank Reactor Optimisation

The production yield of a 20,000 litre STR with twin Rushton impellers, as analysed in Humbird [2020] and illustrated in Figure 5, is optimised for variations in the diameter of the impeller, D, the axial extent - or width - of the blades, W, the impeller RPM and the gassing rate. The radial extent - or length - of the blades, L is maintained at $D/4$. The impeller hub diameter is maintained to coincide with the radial midpoint of the blade width. The sparged gas has 21% oxygen.

3.1.1 Optimisation Approach 1

System level model A is used with an M-Star CFD characterisation model to optimise for yield, driven by the Ax platform. Table 5 shows the key input parameter values and ranges for three different optimisations. In each optimisation an initial set of simulations are performed using a Sobol technique to determine parameter values, with the parameter values for the remaining simulation determined using the Gaussian Process Expected Improvement (GPEI) acquisition function. The number of Sobol simulations and the choice of acquisition function are implemented as defaults by the Ax platform.

ID	N	ϕ [m]	ε_M [W/kg]	RPM	Q [vvm]	D [m]	W [m]
20-A1	30	2.10^{-3}	2	20 - 120	0.01 - 0.1	0.7 - 1.2	0.1 - 0.3
20-A2	50	2.10^{-3}	100	20 - 120	0.01 - 0.1	0.7 - 1.2	0.1 - 0.3
20-A3	30	$1.5.10^{-3}$	2	20 - 120	0.01 -0.1	0.7 - 1.2	0.1 - 0.3

Table 5: Conditions and ranges of input parameters for optimisations of a 20,000 litre STR using optimisation approach 1. N is the number of simulations, ϕ the bubble diameter at standard temperature and pressure conditions, ε_M the maximum time averaged eddy dissipation rate, RPM the impeller rotation rate, Q the sparging rate, D the impeller diameter and W the width of the impeller blades.

The first optimisation, 20-A1, uses a fixed bubble size distribution of 2mm at standard atmosphere and pressure, a limit of a maximum time averaged turbulence energy dissipation rate of 2 W/kg, which corresponds to the Kolmogorov length scale of 20 μm used by Humbird [2020]. The second optimisation, 20-A2, uses a significantly higher limit of the maximum time averaged turbulence energy dissipation rate, 100 W/kg, taken to represent the upper limit of what CHO cells may be able to tolerate (Chalmers [2015]) which may also be representative of what could be achieved with cell line engineering to increase robustness to hydrodynamic stress. The third optimisation, 20-A3, uses a smaller fixed bubble size distribution compared to 20-A1, keeping all other inputs the same. Table 6 shows the optimised yield and parameter values for each of these optimisations.

ID	Yield [g/L wet]	RPM	Q [vvm]	D [m]	W [m]
20-A1	56.6	40.2	0.0495	0.8299	0.1589
20-A2	101.9	120.0	0.0573	0.8654	0.2771
20-A3	94.5	42.2	0.0534	0.8319	0.3000

Table 6: Predicted optimum yield and corresponding input parameter values from the different optimisations of a 20,000 litre STR using approach 1.

The optimum yield corresponds to that achievable from a bioreactor performance perspective, not including metabolic product concentration limits. Optimisation 20-A1 compares directly to the predicted optimum from Humbird [2020], with the only differences being in this case the mass transfer, mixing time and impeller induced hydrodynamic stress being computed on a per configuration basis, rather than from empirical equations, which provides a more robust approach to assess variations in impeller geometry and RPM. Doing so increases the predicted optimum yield by nearly 25% from 45.4 to 56.6 grams per litre wet.

Optimisation 20-A2 considers the effect of a much higher tolerance to induced stress, which is reflected in the increase in optimum yield to 101.9 grams per litre wet. The results indicate that even higher yields may be possible with such a high tolerance, as the optimum impeller RPM is at the upper limit (120) of the input range used, indicating that higher mixing and mass transfer may be achievable at higher RPM due to increased stress tolerance.

Optimisation 20-A3 considers the effect of a smaller initial bubble size, representative of a sparger with a smaller pore size. In such a case the optimum yield is close to that from the much increased stress limit, 94.5 grams per litre wet. It is of interest to note that in 20-A3 the optimum yield is achieved at a higher RPM than for 20-A1, which suggests that the 20-A1 optimisation may not have found the true optimum. This is most likely due to the set-up of the Bayesian optimisation in the default Ax workflow (the acquisition function) and the implemented limit of $N = 30$ simulations. Figure 7 shows the evolution of the predicted yield with iteration N for optimisation 20-A1. The optimum yield is found halfway through the 30 iterations, after which no improvements were found, which is suggestive of a sub-optimal set-up of the optimisation approach.

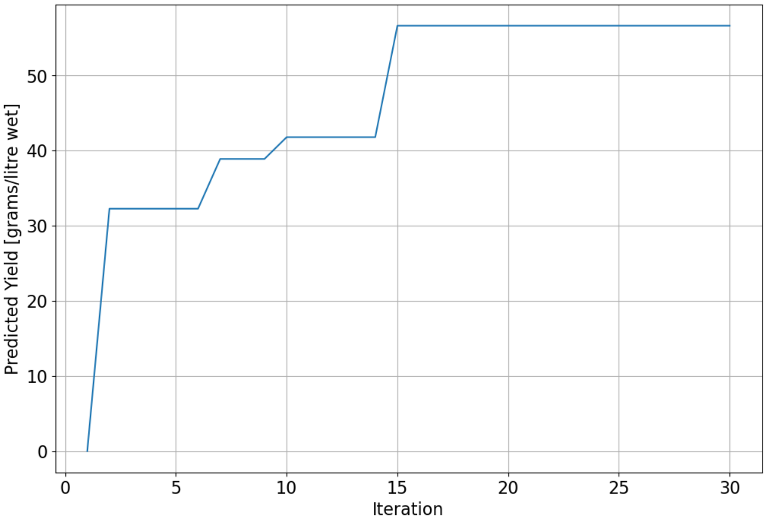

Figure 7: Evolution of optimum yield with loop iteration for optimisation 20-A1.

3.1.2 Optimisation Approach 2

In the second optimisation approach, system level model B is used in place of model A, and a single optimisation is performed with equivalent inputs and Ax model setup to the first optimisation using approach 1 with the higher cell stress limit (20-A2), summarised in Table 7. Table 8 gives the predicted optimum yield and the corresponding input parameter values.

ID	N	ϕ [m]	ε_M [W/kg]	RPM	Q [vvm]	D [m]	W [m]
20-B	30	2.10^{-3}	100	20 - 120	0.01 - 0.1	0.7 - 1.2	0.1 - 0.3

Table 7: Conditions and ranges of input parameters for a 20,000 litre STR using optimisation approach 2. N is the number of simulations, ϕ the bubble diameter at standard temperature and pressure conditions, ε_M the maximum time averaged eddy dissipation rate, RPM the impeller rotation rate, Q the sparging rate, D the impeller diameter and H the height of the impeller blades.

ID	Yield [g/L wet]	RPM	Q [vvm]	D [m]	W [m]
20-B	147.2	94.2	0.1	0.8126	0.2496

Table 8: Predicted optimum yield and corresponding input parameter values from the optimisation of a 20,000 litre STR using approach 2.

Using system level model B, which uses a more robust approach to predict the carbon dioxide concentration limit but does not consider a limit for the sparging rate, predicts a much higher optimum yield (147.2 g/L wet) than using system level model A (101.9 g/L wet). The optimum parameter values show that with system level model B an optimum has been returned at the limit of the sparging rate, whereas system level model A returned an optimum at the limit of impeller RPM.

3.2 42,000 Stirred Tank Reactor Optimisation

A 42,000 litre STR is optimised using system level model B, with inputs shown in Table 9 and predicted optimum yield and corresponding parameters in Table 10.

ID	N	ϕ [m]	ε_M [W/kg]	RPM	Q [vvm]	D [m]	W [m]
42-B	30	2.10^{-3}	100	40 - 160	0.01 - 0.1	0.9 - 1.5	0.15 - 0.4

Table 9: Conditions and ranges of input parameters for a 42,000 litre STR using optimisation approach 2. N is the number of simulations, ϕ the bubble diameter at standard temperature and pressure conditions, ε_M the maximum time averaged eddy dissipation rate, RPM the impeller rotation rate, Q the sparging rate, D the impeller diameter and H the height of the impeller blades.

ID	Yield [g/L wet]	RPM	Q [vvm]	D [m]	W [m]
42-B	131.2	74.3	0.083	0.998	0.323

Table 10: Predicted optimum yield and corresponding input parameter values from the optimisation of a 42,000 litre STR using approach 2.

The predicted optimum yield for a 42,000 litre STR (131.2 g/L wet) is significantly higher than the equivalent from Humbird [2020] (43.5 g/L wet for 21% oxygen sparge gas) and the value used by Negulescu et al. [2023] (100 g/L wet). The optimum parameter values show that neither limits of RPM or sparging rate were reached, although the larger impeller diameter and RPM may have resulted in the eddy dissipation rate limit being reached.

3.3 Stirred Tank Reactor Spatial Variability

The optimum conditions for the 20,000 and 42,000 litre STRs from optimisations 20-B and 42-B are used in a CFD model to predict the effects of spatial variability of dissolved oxygen concentration. The models are each run for 250 and 180 seconds of bioreactor operation respectively, with the minimum, maximum and mean oxygen consumption constraint, $C_{O_2}/(C_{O_2}+K_m)$, from equation 27 becoming independent of time after 120 seconds in both cases. In both cases the model predicted no restrictions to cell growth due to oxygen restriction. Figure 8 shows the distribution of dissolved oxygen at the end of the operational duration for each bioreactor, with both bioreactors showing little variation around a mean concentration of around 40 to 50% of the saturated dissolved oxygen concentration ($\approx 3 \times 10^{-4}$ moles per litre).

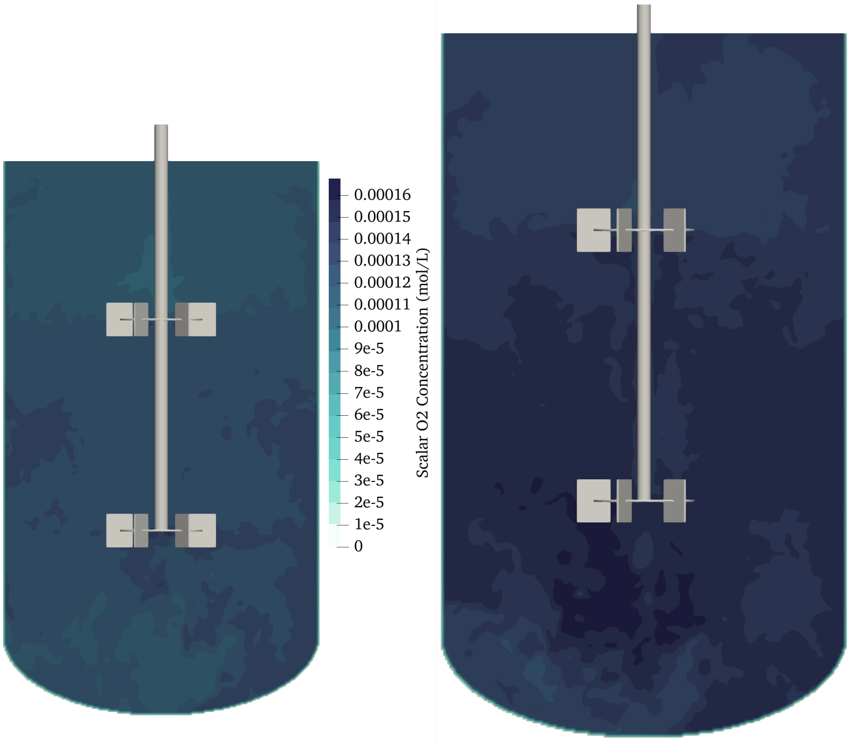

Figure 8: Mid plane contour plot of the dissolved oxygen concentration in 20,000 litre (left) and 42,000 litre (right) STRs after 250 and 180 seconds of operation respectively.

3.4 300,000 Air Lift Reactor Performance

A large 262,000 litre ALR was considered in the TEA of Negulescu et al. [2023] where a maximum cell density of 100 grams per litre wet was used, and modelled in the work of Li et al. [2020] through a combination of 3D computational fluid dynamics for the flow field and dispersed bubble concentration and a 1D advection-reaction-dispersion equation for the dissolved oxygen concentration, where they conclude that a cell density of 2.56×10^8 cells per millilitre, or 770 grams per litre wet, is viable from an oxygen transfer perspective, despite the cell density used being twice that cited from other studies (Van der Weele and Tramper [2014]). Li et al. [2020] use a oxygen uptake rate independent of dissolved oxygen concentration, an order or magnitude lower than the rate used by Humbird [2020] and nearly an order of magnitude lower than the peak uptake rate $q_{\max}$ from Meuwly et al. [2006] used in this study (see Figure 6).

Given the large difference in achievable cell density for a bioreactor of common architecture and scale, the spatial variability CFD model is used to assess the viability of the cell densities used in these studies. The CFD model uses the geometry from Li et al. [2020], including the baffles at the top of the riser and the porous regions intermittently spaced in the riser which were both found to improve performance, illustrated in Figure 9. The gassing rate is the same as used by Li et al. [2020], whereas the bubble diameter at standard temperature and pressure is 2mm here compared to ≈1.6mm from Li et al. [2020]. The spatial resolution is much higher than used by Li et al. [2020] with 70 lattice elements across the down-comer diameter in the present study compared to 16 finite volumes in Li et al. [2020], facilitated by the higher computational efficiency of the lattice Boltzmann approach in M-Star compared to the finite volume approach. Two simulations are performed, one using the cell density from Negulescu et al. [2023], the other using the cell density from Van der Weele and Tramper [2014].

Figure 10 shows contours of the oxygen uptake rate constraint, $C_{O_2}/(C_{O_2}+K_m)$, on a cut plane through the centre of the bioreactor. For optimum performance the constraint will be one everywhere, indicating that the cells are growing at the maximum rate everywhere. The model with the cell density of Van der Weele and Tramper [2014], shown left, shows large areas of no oxygen consumption, corresponding to areas of low availability, indicating that such a cell density is not viable. The image right shows the result from a model with the cell density used by Negulescu et al. [2023], showing small regions of limited uptake. Given the limited size of these regions, it is reasonable to expect that small changes in geometry and/or operating conditions would enable this cell density to be viable.

Figure 9: Schematic of the 300,000 litre airlift reactor with probe locations (purple), sparger (cyan) and media height (red line).

Figure 10: Mid plane contour plot of the oxygen uptake constraint from the spatial variability CFD model of the 300,000 litre ALR. Cell density from Li et al. [2020] on the left, from Negulescu et al. [2023] on the right

4 Discussion

4.1 Bioreactor Performance Modelling

The workflows developed and used here have enhanced prior modelling approaches by enabling the prediction of bioreactor performance characteristics specific to the geometry and operation, as opposed to from empirical equations that are constrained by the parameters in the data set used to generate them. The system level performance models can take such characteristics from alternate sources, which may be measurements or other models, enabling a broader application. The use of CFD to provide specific performance characteristics is beneficial in enabling the full variety of geometry and operation to be considered. The challenge now is the accuracy of these CFD models, which themselves rely on model equations for lower level characteristics. Oxygen mass transfer is one such case where a previously used empirical equation for the system level mass transfer has been replaced with a CFD model which implements an empirical equation for the mass transfer on a per bubble parcel basis. There is little scope for removing empirical equations at this scale as capturing the dynamics of bubble gas liquid interfaces and the resulting mass transfer at all scales is computationally intractable for applications with very large numbers of bubbles as is the case here. The validation of these empirical equations and the CFD model as a whole for the bioreactor(s) under consideration is crucial to the practical application of this approach.

The simple coupling to cell oxygen metabolism has enabled the first step in incorporating the effects of spatial inhomogeneity in to workflows that aim to optimise for production yield. The challenge with such models is the requirement to assess the viability of a cell density rather than optimising for the maximum achievable cell density. The coupling could be easily extended to consider production of carbon dioxide and other metabolites, such as lactate and ammonia as considered previously (Humbird [2020]), the challenge being the availability of data or models that provide the production rates given the local nutrient and metabolite concentrations for the cell line concerned. The creation of such data for cultivated meat cell lines, rather than the common approach of using CHO cells due to their extensive use in research and application, would enable the extension of the modelling approach used here to consider more nutrients and metabolites. The extension of the CFD model to then consider alternate modes of production such as perfusion would be relatively straight forward.

4.2 Optimisation Technique and Computational Efficiency

The Ax platform used here provided a straight forward and user friendly implementation of Bayesian optimisation. However, the approach itself was likely sub-optimal in terms of the efficiency of the automatically selected acquisition function in finding the optimum yield. More specific crafting of the optimisation aspect of the workflow may enable more efficient determination of the optimum yield, and indeed may further increase the predicted optimum yield. Such crafting may consider the use of more than one acquisition function, where one that is better at exploring the parameter space is used initially followed by a second that is better at identifying optima, rather than the use of the GPEI acquisition function, which aims to balance parameter space *exploration* and *exploitation*. The present work suggests that GPEI struggles to balance these effectively in the current application. Approaches that are superior to that employed here may already be reported in the literature or could be developed with moderate effort.

The duration required to solve the CFD models is critical to enabling the optimisation to complete in meaningful times. The use of lattice Boltzmann CFD software on GPUs was critical to enabling this. The STR CFD models used in the optimisations took around three hours to solve on twin NVidia RTX A6000 GPUs, with the $N = 30$ and $N = 50$ run optimisations therefore taking around ninety and one hundred and fifty hours respectively. Running on more GPUs could reduce this duration through computational efficiency and improvement to the optimisation workflow may further reduce the duration through being able to find optima with less parameter space exploration (lower N).

4.3 Bioreactor Performance Potential

The results of the bioreactor optimisations and performance assessments indicate that the maximum achievable cell density in prior TEAs may be lower than is achievable. When comparing to an equivalent 20,000 litre fed batch STR under equivalent constraints the relatively simple optimisation of Rushton impeller blade geometry suggested an improvement of nearly 25% over previously used maximum cell densities in the TEA of Humbird [2020]. Further optimisations suggest that the maximum achievable density may be much higher, through either modification of the bubble diameter through sparger specification or the development or more thorough understanding of cell response to hydrodynamic stress. With a higher tolerance to hydrodynamic stress the results for a 42,000 litre fed batch STR maximum cell density show an increase of over 200% compared to the value derived and used in Humbird [2020], and an increase of over 30% compared to the value used in Negulescu et al. [2023]. The CFD model predicting the effects of spatial variability on oxygen consumption, and by proxy cell viability, suggests that the predicted optimum yields at both 20,000 and 42,000 litres are unlikely to affected by such variations.

The assessment of a 300,000 litre ALR suggests that the cell density used in a prior TEA (100 grams per litre wet Negulescu et al. [2023]) is likely to be viable with small modifications geometry or operation. The assessment also suggests that the high cell density considered viable from prior CFD modelling (770 grams per litre wet Li et al. [2020]) is unlikely to be viable. The uncertainty here is due to the difference in bubble diameter, which was shown to have a significant effect on yield in the 20,000 litre STR optimisations in the present study. Those results showed a near 100% increase in yield when moving from 2mm to 1.5mm bubble diameters, which if consistent across to the 300,000 litre ALR would increase the yield to 200 grams per lite wet. Given the significantly larger volume of the ALR (a factor of 15) one would expect the effect of the reduction in bubble size on yield to be greater than the 100% seen in the 20,000 litre STR. In order to sustain a cell density of 770 grams per lite wet however, an increase of 770% would be required, which may not be achievable even at the much larger volume. The order of magnitude difference in the oxygen uptake rate (Figure 6) used in Li et al. [2020] and the present study is likely to be the main driver for the differences in predicted viable cell densities. The ALR geometry considered here was developed with a few specific modifications by Li et al. [2020] (porous plates and baffles at the riser/down-comer transition) and further optimisation using an approach similar that used here may show significant improvements to the 100 grams per litre wet viable performance shown here.

It is important to recognise that the bioreactor performance potential considered here has focused solely on the performance characteristics of the bioreactors in fed batch mode, and has not considered cell density limitations of concentrations of metabolic products. The reason for this separation was to illustrate what may be achievable from a production hardware perspective and to view this separately to what may be achievable from a cell metabolism and metabolic concentration response perspective. The inclusion of cell density limits due to metabolite concentrations such as lactate and ammonia are straight forward to implement in both the system level and CFD performance models used here. Again, knowledge of such characteristics of cultivated meat cell lines are *critical* in using these limits effectively.

4.4 Recommendations

Techno-economic analyses outputs of cost of goods sold for cultivated meat products are influential in driving investment and production facility development, placing importance on the accuracy of the production yield used or derived in the analyses. The present work suggests achievable yields may be significantly higher than those used in prior TEA's; a simple next step would be to re-evaluate predicted cost of goods using the predicted maximum yields from this study in existing TEA workflows.

A sensible subsequent development would be to include the economic analysis in to the optimisation workflow, allowing optimisation directly for the cost of good sold. Such an approach would then consider both the production performance and operating cost of impellers and spargers, although the operating cost (kWh of power) from driving the impeller and sparger is expected to make a very small contribution to the total operating cost component of the cost of goods sold (Humbird [2020]). A more adventurous development could consider the inclusion of representative cell metabolisms and media amino acid composition, which makes up the majority, $\approx 80\%$, of the fed batch operating costs (Humbird [2020]).

The yields predicted from this study have been evaluated from a bioreactor mixing, mass transfer, hydrodynamic stress and gas hold up performance perspective only, without consideration of the management of metabolic products and the consequential limits to yield. This approach was taken for two reasons; to focus solely on bioreactor performance to ascertain what may be achievable from a hardware design and operation perspective, and to remove the uncertainty in using metabolic behaviour that may not be representative of the animal cell line under consideration.
The use of CHO cell metabolisms is common in published cultivated meat TEAs due to the lack of available data on cell lines being developed in industry. This is a significant source of uncertainty in the predicted achievable yield both from the potential difference in metabolism (primarily nutrient uptake and production rates) between the modelled metabolism and that of the cell line under consideration, and in the variability that exists in modelled metabolisms in previous studies. Figure 6 highlights this clearly, showing an order of magnitude difference in oxygen uptake rates used in different previous studies, in both cases independent of dissolved oxygen concentration, compared to the model used in the present study which is close to the higher value of previous studies but dependent on concentration. This difference is likely a primary cause of the variation in predicted achievable yields from the previous and present studies. The use of CHO cell metabolisms also precludes the ability of models to predict which out of a range of cultivated meat cells lines (pork, beef, chicken, salmon, etc) may be the most cost effective to produce as a result of having the least challenging metabolism to manage from a production perspective.
A similar situation exists with regards the ability of a cell line to grow unaffected by hydrodynamic stress, where CHO cell line data is again used in models, with significant variability in published limiting values. Noticeable simplifications are made with regards bubble induced stress, with limits to sparging rate in vessel volumes per minute or superficial velocity used previously, without consideration of the bubble diameter which directly affects the energy dissipation rate during bubble collapse at the headspace gas-liquid free surface. The 20,000 litre optimisations performed here show the sensitivity of the predicted achievable yield to cell stress response, with a near 100% increase in yield achieved through varying stress tolerance within published.
The use of metabolic and stress response behaviour of the cell line under consideration is therefore *critical* in enabling technical models (TEA, CFD, optimisation and coupled models) to add value to the industry by enabling accurate production facility design and optimisation and financial de-risking of production scale capital investment. Efforts to characterise and share such cell response for use in models would significantly enhance their ability to make meaningful contribution to the success of cultivated meat, both from an accuracy perspective and from the subsequent ability of models (primarily CFD) to be developed to assess production scenarios outside of fed-batch, e.g. perfusion, that may be superior in cost of good sold performance.

The work required to understand and utilise the cellular response to hydrodynamic stress, from both impeller agitation and bubble motion and collapse is considerable and complex. Prior work has involved exposing cells to controlled energy dissipation rates Ma et al. [2002] and measuring the effects on viability, which could offer insight when combined with cell stress experience predicted directly from CFD. These CFD predictions could be achieved by extending the ”lifeline” approach used in CFD models to capture microbe metabolic response to nutrient concentration (Blöbaum et al. [2023]; Haringa et al. [2016]; Haringa [2022]). Such extension would be challenging with regards the accuracy of the stress experienced in the presence of the impeller. Figure 3 (bottom row, middle graph) shows that even at high spatial resolutions ($N_X = 800$), which would require significant increases in computational hardware and time to run the optimisations in this study than used here, the predictions of energy dissipation rate in the impeller discharge region show significant difference to experiment. It should also be noted that experimental measurements of energy dissipation rate around impeller are without significant challenge themselves. Modelling of cell stress lifelines in the presence of bubble motion and collapse would require CFD models that captured the evolution of the bubble interface and the resulting energy dissipation in the liquid across a range of scenarios, including but not limited to bubble velocities, bubble sizes, bubble collapse conditions and surfactant concentrations. Given these challenges to understanding the generation of hydrodynamic stress and the effects on cell metabolism, designing novel bioreactors that can cultivate cells without exposing them to impellers or bubbles could be a more cost and time effective route.

The optimisation workflows developed use a simple front end (the Ax platform) to the BOTorch library, which facilitates ease of use but encounters difficulty balancing exploration and exploitation of the parameter space. An assessment of other acquisition functions, a deeper dive in to the characteristics of the yield response in the considered parameter space, or a more detailed implementation in BOTorch directly may provide improvements to the optisation workflow and may also reveal higher potential optimum yields. Such approaches could then be further developed to consider a wider range of parameters than considered here, such as bubble size distributions rather than single values, and the wide range of potential geometric optimisations including but not limited to the broad range of impeller designs, axial impeller location, axial location of one or more spargers and the geometry specific to air lift reactors such as porous plate and baffles.

5 Conclusions

Open access workflows have been developed and used to predict optimum cultivated meat yield in stirred tank and airlift bioreactors. The workflows use and build on prior approaches and couple to computational fluid dynamics models and a Bayesian optimisation routine. A metabolic model of oxygen uptake is included in a computational fluid dynamics model to provide predictions of the effects of spatial in-homogeneity of dissolved oxygen concentration, which is used to indicate if such effects may impact the predicted optimum yield. Optimisations of 20,000 and 42,000 litre fed batch stirred tank reactors suggest that achievable yield may be significantly higher than that used in previous techno-economic analyses, from a bioreactor performance perspective (neglecting metabolic product concentration constraints). Assessment of a 300,000 litre airlift bioreactor suggests that the yield used in a recent TEA is achievable, but that from a prior computational fluid dynamics analysis may not be. The developed workflows can be easily extended to consider cell metabolism and alternate production modes. The use of cultivated meat cell line characteristics, rather than commonly used CHO characteristics as used here, will be critical in being able to accurately predict and optimise production yield. Further development of the optimisation routine to include capital and operating production costs could result in a valuable tool in enabling cultivated meat to overcome production cost challenges.

6 Acknowledgements

This work was funded by the Good Food Institute. Thanks to Pieter Brorens at TU Delft for technical feedback on the first version.

References

Luisa Blöbaum, Cees Haringa, and Alexander Grünberger. Microbial lifelines in bioprocesses: From concept to application. *Biotechnology advances*, 62:108071, 2023.

Jeffrey J Chalmers. Mixing, aeration and cell damage, 30+ years later: what we learned, how it affected the cell culture industry and what we would like to know more about. *Current Opinion in Chemical Engineering*, 10:94–102, 2015.

Jeffrey J Chalmers. The challenges of hydrodynamic forces on cells used in cell manufacturing and therapy. *Current Opinion in Biomedical Engineering*, 20:100357, 2021.

MFW Distelhoff, J Laker, AJ Marquis, and JM Nouri. The application of a strain gauge technique to the measurement of the power characteristics of five impellers. *Experiments in fluids*, 20(1):56–58, 1995.

A Ducci and M Yianneskis. Direct determination of energy dissipation in stirred vessels with two-point lda. *AIChE journal*, 51(8):2133–2149, 2005.

FAO. Emissions due to agriculture. global, regional and country trends 2000–2018. 2021.

Hooman Yadollahi Farsani, Johannes Wutz, Brian DeVincentis, John A Thomas, and Seyed Pouria Motevalian. Modeling mass transfer in stirred microbioreactors. *Chemical Engineering Science*, 248:117146, 2022.

Jürgen Fitschen, Sebastian Hofmann, Johannes Wutz, AV Kameke, Marko Hoffmann, Thomas Wucherpfennig, and Michael Schlüter. Novel evaluation method to determine the local mixing time distribution in stirred tank reactors. *Chemical engineering science: X*, 10:100098, 2021.

Felix Garcia-Ochoa, Emilio Gomez, Victoria E Santos, and Jose C Merchuk. Oxygen uptake rate in microbial processes: an overview. *Biochemical engineering journal*, 49(3):289–307, 2010.

Greg L Garrison, Jon T Biermacher, and B Wade Brorsen. How much will large-scale production of cell-cultured meat cost? *Journal of Agriculture and Food Research*, 10:100358, 2022.

David R Gray, Su Chen, William Howarth, Duane Inlow, and Brian L Maiorella. Co 2 in large-scale and high-density cho cell perfusion culture. *Cytotechnology*, 22:65–78, 1996.

Cees Haringa. An analysis of organism lifelines in an industrial bioreactor using lattice-boltzmann cfd. *Engineering in Life Sciences*, 23(1):1–16, 2022.

Cees Haringa, Wenjun Tang, Amit T Deshmukh, Jianye Xia, Matthias Reuss, Joseph J Heijnen, Robert F Mudde, and Henk J Noorman. Euler-lagrange computational fluid dynamics for (bio) reactor scale down: an analysis of organism lifelines. *Engineering in life sciences*, 16(7):652–663, 2016.

Cees Haringa, Ruben Vandewijer, and Robert F Mudde. Inter-compartment interaction in multi-impeller mixing: Part i. experiments and multiple reference frame cfd. *Chemical Engineering Research and Design*, 136:870–885, 2018.

David Humbird. Scale-up economics for cultured meat: Techno-economic analysis and due diligence. *engRxiv*, pages 2020–12, 2020.

David Humbird. Scale-up economics for cultured meat. *Biotechnology and Bioengineering*, 118(8):3239–3250, 2021.

M Jahoda, M Moštěk, A Kukuková, and V Machoň. Cfd modelling of liquid homogenization in stirred tanks with one and two impellers using large eddy simulation. *Chemical Engineering Research and Design*, 85(5):616–625, 2007.

Stephan C Kaiser, Sören Werner, Valentin Jossen, Matthias Kraume, and Dieter Eibl. Development of a method for reliable power input measurements in conventional and single-use stirred bioreactors at laboratory scale. *Engineering in life sciences*, 17(5):500–511, 2017.

Athanasios G Kanaris, Theodosios I Pavlidis, Ariadni P Chatzidafni, and Aikaterini A Mouza. The effects of the properties of gases on the design of bubble columns equipped with a fine pore sparger. *ChemEngineering*, 2(1):11, 2018.

Y Kawase and M Moo-Young. Mathematical models for design of bioreactors: Applications of: Kolmogoroff's theory of isotropic turbulence. *The Chemical Engineering Journal*, 43(1):B19–B41, 1990.

Minsu Kim, Sunghyun Cho, Areum Han, Yoojin Han, Joseph Sang-Il Kwon, Jonggeol Na, and Il Moon. Multi-objective bayesian optimization for design and operating of fluidized bed reactor. In *Computer Aided Chemical Engineering*, volume 49, pages 1297–1302. Elsevier, 2022.

Maike Kuschel, Jürgen Fitschen, Marko Hoffmann, Alexandra von Kameke, Michael Schlüter, and Thomas Wucherpfennig. Validation of novel lattice boltzmann large eddy simulations (lb les) for equipment characterization in biopharma. *Processes*, 9(6):950, 2021.

Xueliang Li, Guoqiang Zhang, Xinrui Zhao, Jingwen Zhou, Guocheng Du, and Jian Chen. A conceptual air-lift reactor design for large scale animal cell cultivation in the context of in vitro meat production. *Chemical Engineering Science*, 211:115269, 2020.

M-Star. M-star documentation. URL `https://docs.mstarcfd.com/`.

Ningning Ma, Kurt W Koelling, and Jeffrey J Chalmers. Fabrication and use of a transient contractional flow device to quantify the sensitivity of mammalian and insect cells to hydrodynamic forces. *Biotechnology and bioengineering*, 80 (4):428–437, 2002.

F Meuwly, F Loviat, P-A Ruffieux, AR Bernard, A Kadouri, and U von Stockar. Oxygen supply for cho cells immobilized on a packed-bed of fibra-cel® disks. *Biotechnology and bioengineering*, 93(4):791–800, 2006.

William M Miller, Charles R Wilke, and Harvey W Blanch. Effects of dissolved oxygen concentration on hybridoma growth and metabolism in continuous culture. *Journal of cellular physiology*, 132(3):524–530, 1987.

BN Murthy and JB Joshi. Assessment of standard k–ε, rsm and les turbulence models in a baffled stirred vessel agitated by various impeller designs. *Chemical engineering science*, 63(22):5468–5495, 2008.

Patrick G Negulescu, Derrick Risner, Edward S Spang, Daniel Sumner, David Block, Somen Nandi, and Karen A McDonald. Techno-economic modeling and assessment of cultivated meat: Impact of production bioreactor scale. *Biotechnology and Bioengineering*, 120(4):1055–1067, 2023.

Alvin W Nienow. Reactor engineering in large scale animal cell culture. *Cytotechnology*, 50(1-3):9, 2006.

Alvin W Nienow. Impeller selection for animal cell culture. *Encyclopedia of Industrial Biotechnology: Bioprocess, Bioseparation, and Cell Technology*, pages 1–25, 2009.

Alvin W Nienow. Stirring and stirred-tank reactors. *Chemie Ingenieur Technik*, 86(12):2063–2074, 2014.

Paris OECD. Oecd-fao agricultural outlook 2021-2030. 2021.

Seongeon Park, Jonggeol Na, Minjun Kim, and Jong Min Lee. Multi-objective bayesian optimization of chemical reactor design using computational fluid dynamics. *Computers & Chemical Engineering*, 119:25–37, 2018.

Stephen B Pope. *Turbulent flows*. Cambridge university press, 2000.

VV Ranade and JB Joshi. Flow generated by a disc turbine. part ii. mathematical modelling and comparison with experimental data. *Chemical Engineering Research and Design*, 68(1):34–50, 1990.

Derrick Risner, Fangzhou Li, Jason S Fell, Sara A Pace, Justin B Siegel, Ilias Tagkopoulos, and Edward S Spang. Preliminary techno-economic assessment of animal cell-based meat. *Foods*, 10(1):3, 2020.

H Ritchie and M Roser. Environmental impacts of food production. 2020. *Our World In Data*, 2021.

Evan M Schlaich, John A Thomas, Lakshmi Kandari, Gabi Tremml, and Anurag Khetan. Experimental and computational characterization of mass transfer in high turndown bioreactors. *Biotechnology Progress*, page e3330, 2023.

John Thomas, Kushal Sinha, Gayathri Shivkumar, Lei Cao, Marina Funck, Sherwin Shang, and Nandkishor K Nere. A cfd digital twin to understand miscible fluid blending. *Aaps Pharmscitech*, 22:1–14, 2021a.

John A Thomas, Xiaoming Liu, Brian DeVincentis, Helen Hua, Grace Yao, Michael C Borys, Kathryn Aron, and Girish Pendse. A mechanistic approach for predicting mass transfer in bioreactors. *Chemical Engineering Science*, 237: 116538, 2021b.

John A Thomas, Brian DeVincentis, Navraj Hanspal, and Richard O Kehn. Predicting gas-liquid mass transfer rates in reactors using a bubble parcel model. *Chemical Engineering Science*, 264:118183, 2022a.

John A Thomas, Brian DeVincentis, Johannes Wutz, and Francesco Ricci. Predicting the diameters of droplets produced in turbulent liquid–liquid dispersion. *AIChE Journal*, 68(7):e17667, 2022b.

John A Thomas, Anisur Rahman, Johannes Wutz, Ying Wang, Brian DeVincentis, Brendan McGuire, and Lei Cao. Modeling free surface gas transfer in agitated lab-scale bioreactors. *Chemical Engineering Communications*, 210(8): 1328–1339, 2023.

Department of Economic United Nations and Population Division (2022) Social Affairs. World population prospects 2022: Methodology of the united nations population estimates and projections (un desa/pop/2022/tr/no. 4), 2022.

Cor Van der Weele and Johannes Tramper. Cultured meat: every village its own factory? *Trends in biotechnology*, 32(6): 294–296, 2014.

Robert Vergeer, Pelle Sinke, and Ingrid Odegard. Tea of cultivated meat future projections for different scenarios - corrigendum. 2021.

John Villadsen, Jens Nielsen, and Gunnar Lidén. *Bioreaction engineering principles*. Springer Science & Business Media, 2011.

H Wu and GK Patterson. Laser-doppler measurements of turbulent-flow parameters in a stirred mixer. *Chemical engineering science*, 44(10):2207–2221, 1989.

H Wu, GK Patterson, and M Van Doorn. Distribution of turbulence energy dissipation rates in a rushton turbine stirred mixer. *Experiments in Fluids*, 8(3-4):153–160, 1989.

Zizhuo Xing, Brian M Kenty, Zheng Jian Li, and Steven S Lee. Scale-up analysis for a cho cell culture process in large-scale bioreactors. *Biotechnology and bioengineering*, 103(4):733–746, 2009.

Luoding Zhu, Guowei He, Shizhao Wang, Laura Miller, Xing Zhang, Qian You, and Shiaofen Fang. An immersed boundary method based on the lattice boltzmann approach in three dimensions, with application. *Computers & Mathematics with Applications*, 61(12):3506–3518, 2011.