

Article

From Data to Draught: Modelling and Predicting Mixed-Culture Beer Fermentation Dynamics Using Autoregressive Recurrent Neural Networks

Alexander O'Brien ¹, Hongwei Zhang ^{1,*}, Daniel M. Allwood ² and Andy Rawsthorne ³

¹ National Centre of Excellence for Food Engineering, Sheffield Hallam University, Howard Street, Sheffield S1 1WB, UK; alexander.o'brien@shu.ac.uk

² Biomolecular Sciences Research Centre, Sheffield Hallam University, Howard Street, Sheffield S1 1WB, UK; d.allwood@shu.ac.uk

³ Department of Engineering and Mathematics, Sheffield Hallam University, Howard Street, Sheffield S1 1WB, UK; a.rawsthorne@shu.ac.uk

* Correspondence: h.zhang@shu.ac.uk

Abstract: The ascendancy of the craft beer movement within the brewing industry may be attributed to its commitment to unique flavours and innovative styles. Mixed-culture fermentation, celebrated for its novel organoleptic profiles, presents a modelling challenge due to its complex microbial dynamics. This study addresses the inherent complexity of modelling mixed-culture beer fermentation while acknowledging the condition monitoring limitations of craft breweries, namely sporadic offline sampling rates and limited available measurement parameters. A data-driven solution is proposed, utilising an Autoregressive Recurrent Neural Network (AR-RNN) to facilitate the production of novel, replicable, mixed-culture fermented beers. This research identifies time from pitch, specific gravity, pH, and fluid temperature as pivotal model parameters that are cost-effective for craft breweries to monitor offline. Notably, the autoregressive RNN fermentation model is generated using high-frequency multivariate data, a departure from intermittent offline measurements. Employing the trained autoregressive RNN framework, we demonstrate its robust forecasting prowess using limited offline input data, emphasising its ability to capture intricate fermentation dynamics. This data-driven approach offers significant advantages, showcasing the model's accuracy across various fermentation configurations. Moreover, tailoring the design to the craft beer market's unique demands significantly enhances the model's practicable predictive capabilities. It empowers nuanced decision-making in real-world mixed-culture beer production. Furthermore, this model lays the groundwork for future studies, highlighting transformative possibilities for cost-effective model-based control systems in the craft beer sector.

Keywords: mixed-culture beer fermentation; craft beer; Autoregressive Recurrent Neural Network; data-driven modelling; fermentation dynamics prediction



Citation: O'Brien, A.; Zhang, H.; Allwood, D.M.; Rawsthorne, A. From Data to Draught: Modelling and Predicting Mixed-Culture Beer Fermentation Dynamics Using Autoregressive Recurrent Neural Networks. *Modelling* **2024**, *5*, 201–222. <https://doi.org/10.3390/modelling5010011>

Academic Editor: Antonio Brasiello

Received: 22 December 2023

Revised: 26 January 2024

Accepted: 5 February 2024

Published: 7 February 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Mixed-culture fermentation is a time-honoured brewing tradition, offering a spectrum of flavours and aromas distinct from its monoculture counterparts. Unlike monoculture fermentation, where a single yeast strain exclusively dominates the process, mixed-culture fermentation involves the cooperative action of various yeast strains and, often, bacteria. This collaborative microbial inoculum contributes to the complexity and uniqueness of the final beer product.

Historically, mixed-culture fermentation has been celebrated in traditional beer styles such as Lambic and Belgian ales, where wild yeasts like *Brettanomyces* [1–7] and *Lactobacillus* bacteria [3–5,8–10] play pivotal roles. However, the recent resurgence in interest goes beyond tradition, with craft breweries exploring novel combinations of yeast and bacteria to craft beers with diverse and exciting organoleptic profiles [4,5,11].

Fermentation modelling has traditionally focused on monoculture processes, where the dynamics are more predictable due to the dominance of a single yeast strain [11–13]. Mechanistic kinetic models, based on differential equations, have been foundational in understanding and predicting monoculture beer fermentation processes [13,14]. However, when it comes to mixed-culture fermentations, the cooperative and competitive interactions between microorganisms introduce complexities that challenge traditional modelling approaches [15,16].

Hybrid models, combining mechanistic and empirical elements, have also been employed to model beer fermentation. Chiefly, the hybrid models used in the brewing domain are kinetic models using parameters based on empirical data [14,17–23]. However, this hybrid approach is inflexible, tuned to a specific process, and does not address the complexities of mixed-culture fermentation [22], while there is a wealth of research on Lambic-style spontaneous fermentations from a brewing science perspective, the need for predictive models tailored explicitly to mixed cultures of ale and lager yeasts is apparent [24].

The resurgence of mixed-culture fermentation in craft brewing necessitates a paradigm shift in modelling approaches. The traditional understanding of fermentation dynamics, rooted in monoculture kinetic models, is insufficient to capture the intricacies of mixed-culture processes. To overcome this, a data-driven approach becomes imperative.

Data-driven models, particularly those leveraging artificial intelligence methods, offer the flexibility to adapt to the complexities of mixed-culture fermentation [22,25,26]. These models can learn from a diverse range of fermentation scenarios, accommodating variations in yeast strains, ratios, and fermentation temperatures. The versatility of data-driven approaches aligns well with the diverse and experimental nature of craft brewing, providing a pathway to model the dynamics of mixed-culture fermentation accurately.

Advanced empirical methods have been applied to modelling monoculture fermentation [27–30]. However, an Autoregressive Recurrent Neural Network (AR-RNN) approach was selected to model mixed-culture beer fermentation. AR-RNNs excel in capturing temporal dependencies, making them well-suited for modelling dynamic processes with intricate time-variant patterns. In the context of mixed-culture fermentation, where the interactions between different microbial species unfold over time, the ability of AR-RNNs to discern and predict these dependencies is crucial.

By focusing on ale and lager yeast mixed-culture fermentations across various ratios and primary fermentation temperatures, we aim to showcase the adaptability and efficacy of AR-RNNs in capturing the nuances of these dynamic processes.

In the Materials and Methods Section, the methodologies employed for data collection, experimentation, and the application of AR-RNNs to model mixed-culture fermentations are detailed. The steps taken to ensure a comprehensive and representative dataset for training and evaluation are outlined, along with the specifics of AR-RNN architecture and hyperparameter tuning.

Presented in the Results and Discussion Section are the outcomes of our AR-RNN modelling approach. The model performance is determined across different ratios of ale [12,31–36] and lager yeasts [24,33,37–39] and primary fermentation temperatures. Graphical representations and statistical analyses highlight the model's ability to capture the dynamic behaviour of mixed-culture fermentations. The Discussion Section delves into the implications of our results, comparing them with existing models and shedding light on the advantages and limitations of the AR-RNN approach. This section explores the insights gained from the model's predictions, providing context for the observed dynamics in mixed-culture fermentations.

In the final section, the key findings of our study are summarised, and their significance in the broader context of mixed-culture fermentation in craft beer production are presented. The practical implications for brewers and the potential for further advancements in data-driven modelling approaches for complex fermentation processes are also discussed.

2. Materials and Methods

2.1. Materials

The experimental data for this study were generated at the National Centre of Excellence for Food Engineering (NCEFE) using our on-site pilot-scale brewhouse facility with a capacity of 200 L. The brewhouse setup comprised a 200 L Speidels Braumeister (Ofterdingen, Germany) (Figure 1) and a 240 L Speidels fermentor (Ofterdingen, Germany) (Figure 2). To enable the real-time monitoring of fermentation parameters, a BrewIQ fermentation monitor (Figure 3), an Internet of Things (IoT) sensor array from Precision Fermentation, was used to sample pilot scale fermentations online.



Figure 1. The 200 L Speidels Braumeister used in the brewing process (image courtesy of Speidels).

The brewing process involved the production of beer using a Speidels Braumeister (Figure 1), following a grain bill for an India Pale Ale (IPA). The grain bill consisted of pale ale malt and Carapils malt, while the hop addition schedule included centennial and cascade hops.



Figure 2. The 240 L Speidels fermentor used in the brewing process (image courtesy of Speidels).



Figure 3. The Precision Fermentation BrewIQ fermentation monitor used for real-time monitoring of fermentation parameters.

The dataset employed in this study involved brewing beer with both monocultures and mixed cultures. Two representative yeast species, *Saccharomyces cerevisiae* (Fermentis SafAle US-05 (ale yeast)) and *Saccharomyces pastorianus* (Fermentis SafLager S-23 (lager yeast)), were used in various ratios. Fermentations were conducted at two different temperatures: 22 degrees Celsius and 12 degrees Celsius. Subsequently, the batches were conditioned at 5 degrees Celsius after a fermentation period of 5 days. Figure 4 showcases an example of the beers produced during the data generation process.



Figure 4. Beers fermented with different yeast configurations at 22 degrees Celsius. The inocula, left to right, are as follows: (a) ale yeast monoculture inoculum (SafAle US-05); (b) ale and lager yeast (SafAle US-05 and SafLager S-23) mixed-inoculum in a 1:1 ratio; (c) lager yeast monoculture inoculum (SafLager S-23).

For each fermentation configuration, spanning both yeast monoculture and mixed-culture scenarios, data were collected for the initial 300 h (12.5 days) of fermentation. While primary fermentation for a typical IPA typically requires only 3–5 days, monitoring for a

longer period allows for conditioning dynamics to be monitored in addition to capturing the dynamics of slower, lower temperature fermentations.

To enhance the dataset, augmentations of the data were performed. The raw data were cleaned and duplicated before introducing pseudo-random Gaussian noise to the data. The maximum amplitude of the noise added to the data was limited to the stated tolerances of the sensors used to produce the original data (see Table 1). This resulted in the creation of ninety time-series fermentation datasets, encompassing over eight fermentation parameters and comprising nine hundred samples (300 h) of data per batch.

Table 1. Table of the sensor array tolerances for the parameters of interest.

Sensor	Tolerance
Specific Gravity Sensor	± 0.2 °P
pH Sensor	± 0.1
Fluid Temperature Sensor	± 1 °P

2.2. Methods

Before initiating the modelling process, the dataset underwent preprocessing, involving the removal of redundant attributes, ensuring that the input data were stripped of superfluous information. The dataset was resampled to 20 min intervals, standardising the measurement frequency. Target attributes for validation were produced by temporally shifting input attributes, laying the groundwork for the autoregressive aspect of the model. Finally, data normalisation ensured uniformity and comparability across the diverse fermentation scenarios considered in the study.

The prolonged nature of brewing processes, often exceeding 300 h per batch, presented a unique challenge in generating sufficient process data for training an RNN model. Addressing this challenge required a thoughtful approach to data augmentation. Gaussian noise, within the tolerances of the sensor equipment, was introduced deliberately during data augmentation. This noise addition augmented the dataset and introduced a level of variability, ensuring that the model was trained to handle real-world fluctuations and uncertainties. The careful consideration of challenges and the introduction of noise contributed to the robustness and real-world applicability of the finalised AR-RNN model.

The core of the Autoregressive Recurrent Neural Network (AR-RNN) lies in the Recurrent Neural Network (RNN) model, which was made more capable through the use of Long Short-Term Memory (LSTM) units. RNNs, specialised for sequential data, were a natural fit for handling the temporal dependencies inherent in fermentation processes. LSTMs, a variant of RNNs, emerged as critical components due to their prowess in learning and retaining information over extended sequences. This characteristic effectively addressed challenges such as the vanishing gradient problem, ensuring the model's capability to capture intricate, long-term dependencies that characterise mixed-culture fermentation across many hours.

Mixed-culture fermentation processes are inherently more complex than monoculture fermentation, demanding a model capable of handling multivariate, nonlinear, time-variant data in a range of fermentation configurations. LSTMs, equipped with memory cell structures and gating mechanisms, emerged as a fitting choice. The AR-RNN's reliance on LSTMs aligned seamlessly with the intricate and evolving nature of fermentation processes, providing a robust foundation for capturing complex dependencies over time.

The recurrent neural network architecture, a critical element in this data-driven modelling approach, was produced, trained, and tuned using TensorFlow and Keras. TensorFlow, an open-source machine learning framework developed by Google, formed the computational backbone, while the Keras library, operating atop TensorFlow, provided a high-level neural networks API. The coupling of these tools ensured, not only the depth of customisation necessary for our complex fermentation model, but also a user-friendly interface conducive to rapid experimentation and iteration.

The structuring of the RNN topology was aimed at achieving an optimal balance between capturing temporal dependencies and preventing overfitting, including dropout and L2 regularisation. It comprised an LSTM layer followed by a dropout layer, another LSTM layer, a dropout layer, a dense layer, and a time-distributed dense layer (Figure 5). The hyperparameter tuning process had a search space of the following hyperparameters: number of LSTM nodes, dropout rates, dense layer nodes, learning rates, batch sizes, epochs, and L2 regularisation. This search space for the tuning process ensured the RNN was finely calibrated to the specifics of mixed-culture fermentation dynamics.

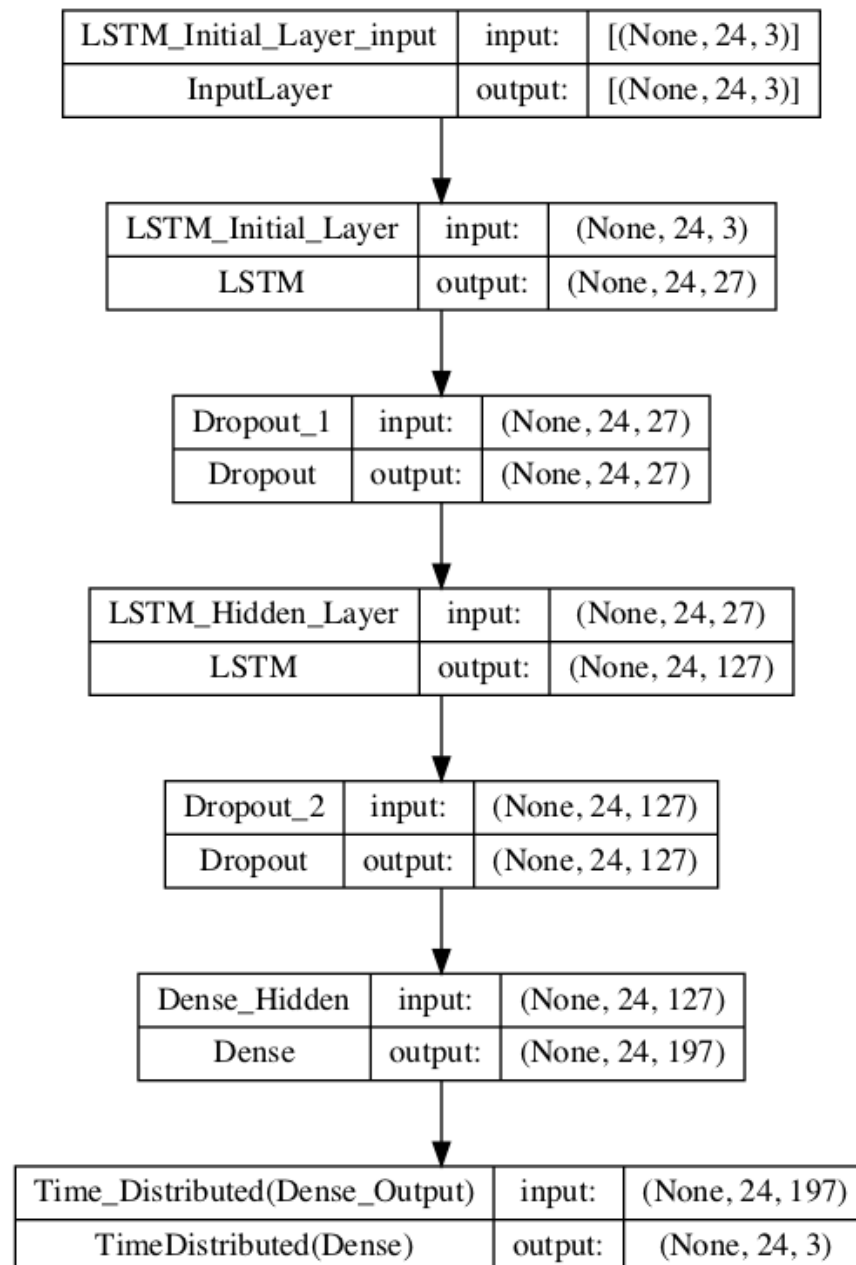


Figure 5. Topology of the RNN model with optimised hyperparameters.

The success of the model hinged on the optimisation of critical hyperparameters. This process involved fine-tuning all hyperparameters within the search space. A Bayesian optimisation approach was adopted to expedite this exploration in the complex parameter space. This method leverages probabilistic models to predict the model's performance, accelerating convergence to optimal configurations. The scikit-optimize Python library

was used to implement this Bayesian optimisation process. Mean Squared Error (MSE), a standard RNN evaluation metric, was selected as the primary evaluation metric for model performance. MSE quantitatively measures the model's predictive accuracy, making it a suitable candidate metric for selecting the best-performing hyperparameter configuration.

Training, testing, and validation data were split in a 70:15:15 ratio, respectively. As the data are sequential batch data, the sequence order of the data were preserved and data were split within their respective batches to ensure agreement between the training, testing, and validation datasets. Furthermore, to ensure the validity of the MSE evaluation metric used during the hyperparameter tuning process, for each call of the optimisation process, each batch was individually assessed for MSE performance. The MSE value used to assess the overall model performance was the mean value of the MSE score across every fold for each optimisation call.

Hyperparameter tuning was essential to strike the right balance between model complexity and overfitting, ensuring the AR-RNN's adaptability to the specifics of monoculture and mixed-culture beer fermentation dynamics.

A distinctive feature of the AR-RNN model lies in autoregression, wherein the input and output dimensions and attributes are the same, enabling the model to predict future values from previous time steps. The model itself takes a series of 24 time steps, representing eight hours of fermentation data, and predicts the subsequent eight hours of fermentation dynamics. The autoregressive mechanism was fundamental in enhancing the model's predictive time horizon capabilities, a critical aspect for capturing the evolving dynamics of mixed-culture fermentation accurately in a real-world environment. The ability to feed predictions back into the model facilitated the generation of forecasts across arbitrary predictive time horizons. This unique capability extended the temporal reach of the model, allowing it to discern and predict complex interactions unfolding over extended durations. The autoregressive aspect of this model, predicting the future values of the measurement attributes exclusively from previous time steps of the same measurement attributes, was a key innovation in this field, elevating the model's capability beyond traditional predictive models that require exogenous data or tuned parameters.

3. Results and Discussion

3.1. Training Data Characteristics

The figures below are corner plots showing example data for each category of inocula at 12 °C and 22 °C for Figures 6 and 7, respectively. The corner plots in Figures 6 and 7 demonstrate the mutual relationships between the key fermentation parameters, namely specific gravity, pH, fluid temperature, and time, in addition to the distributions of each parameter along the leading diagonal.

Figure 6 shows a slow decrease in specific gravity over time, and all inocula configurations fail to attenuate during the 300 h time period fully. Moreover, the differences in the attenuation behaviour between the various inocula are pronounced, whereas the pH dynamics are broadly similar in character, indicating that specific gravity is chiefly affected by the lower primary fermentation temperature of 12 °C.

Conversely, Figure 7 shows the corner plot for attribute data with a primary fermentation temperature of 22 °C. As can be observed, the majority of the inocula combinations manages to fully attenuate at this higher primary fermentation temperature of 22 °C within the 300 h time span. In this instance, the specific gravity data appear to be far more closely mirrored, whereas there is a comparatively higher variance of pH values than the beers fermented at the lower primary fermentation temperature (12 °C).

These data insights support the notion that a data-driven approach is required, as the dynamics for each of the ten configurations shown (five inocula ratios and two temperatures) in Figures 6 and 7 vary significantly and cannot be easily modelled from first principles.

Pairplot of Fermentation Data by Category




Figure 6. Corner plot of a reduced dataset for beer fermentations with a primary fermentation temperature of 12 °C.

3.2. Hyperparameter Tuning

As previously detailed in the Methods Section (Section 2.2), the RNN layer topology was designed to ensure that temporal dependencies were captured within the LSTM layers, overfitting was mitigated with the dropout layers before being passed through a dense and time-distributed layer to produce an output (see Figure 5), the number of layers were fixed as a deeper network may lead to overfitting without a sufficiently large training dataset. However, to produce an optimised model, critical hyperparameters must be tuned, namely LSTM units layer 1, LSTM units layer 2, dropout rate, dense units, learning rate, batch size, epochs, and L2 regularisation.

The hyperparameters were tuned using a Bayesian optimisation approach using the Adam algorithm and the scikit-optimize library. Figures 8 and 9 show the training vs. validation loss plot of the RNN model produced from 25 and 20 calls of the optimisation function, respectively. Figure 8, tuned over 25 iterations, shows an excellent model performance for training data. However, the validation loss trace does not agree with the training loss trace and diverges over time; this behaviour is indicative of significant overfitting.

Pairplot of Fermentation Data by Category

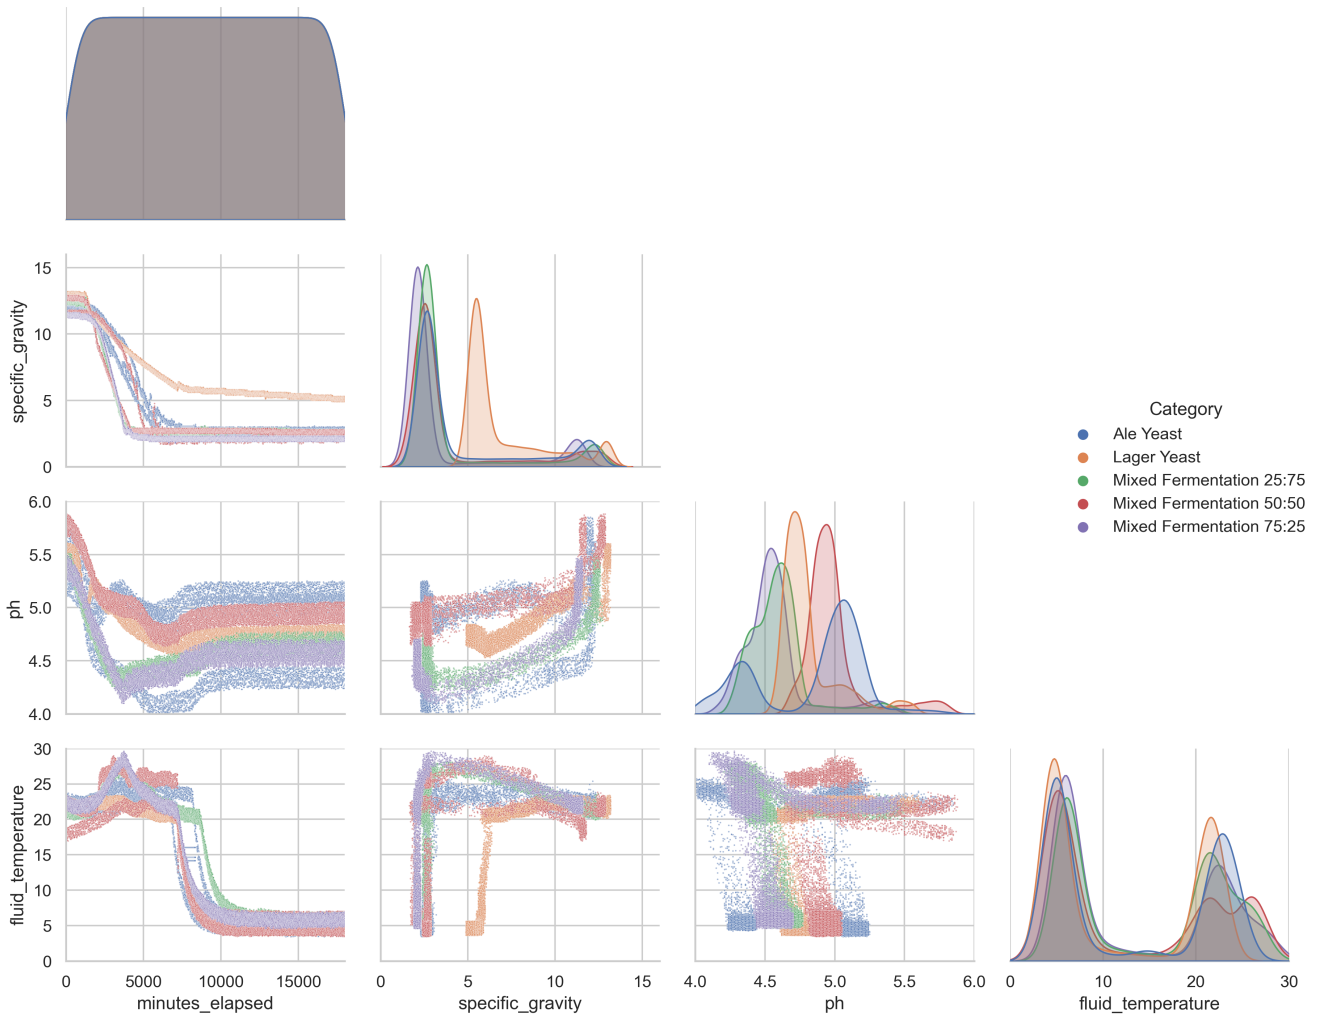


Figure 7. Corner plot of a reduced dataset for beer fermentations with a primary fermentation temperature of 22 °C.




Figure 8. Training vs. validation loss plot for the model produced from 25 calls of the Bayesian optimisation process.




Figure 9. Training vs. validation loss plot for the model produced from 20 calls of the Bayesian optimisation process.

Alternatively, a model produced using only 20 optimisation steps (Figure 9) shows a training vs. validation loss plot against MSE over 71 epochs, demonstrating a good performance, wherein the validation loss trace closely mirrors the training loss trace, indicating that the model is appropriately fitted. Moreover, while the final MSE value for training loss is slightly higher than the model optimised for a further five optimisation steps (Figure 8), the final MSE values for both training and validation loss are comfortably below 0.01 in Figure 9. The complete set of optimised hyperparameters is detailed in Table 2, the impact of which is also reflected within the node values for the model's topology in Figure 5.

Table 2. Table of the tuned hyperparameters produced from 20 Bayesian optimisation steps.

Hyperparameter	Tuned Value
LSTM Units Layer 1	27
LSTM Units Layer 2	127
Dense Units	197
Batch Size	110
Epochs	71
Learning Rate	0.000251355666177083
Dropout Rate	0.2529847965068651
L2 Regularisation	0.000022464551680532603

3.3. Residual Results for the AR-RNN Model

The AR-RNN model is designed to take the previous eight hours (24 samples) to predict the following eight hours of data. Direct feedback into the model allows this time horizon to be extended further. However, errors will compound and cause the predictions to become less accurate over extended predictive horizons.

Figures 10–24 contain example model performance data in the form of residuals for each of the five inocula configurations: monoculture ale yeast, monoculture lager yeast, mixed-culture ale and lager yeast (25:75), mixed-culture ale and lager yeast (50:50), and mixed-culture ale and lager yeast (75:25), respectively. All of the plots are of residual data. The residual plots were produced by splitting unseen datasets into rolling sequences that were 24 time steps long. The predictions made by the model were compared to the real-world, ground-truth data, and the residual (error) values were plotted for each attribute. The green dashed lines indicate the tolerances of the sensor equipment used to obtain the training, testing, and validation datasets listed in Table 1. The red trace is a locally weighted

scatter plot smooth, or lowess line. In the context of these residual plots, the lowess should approximate zero error across time to indicate a good model performance. Moreover, the scatter plot of residual values should be as low as possible and the appearance of noisy, discernible patterns within the residuals scatter plot may indicate that the model has not fully captured or generalised the dynamics of the fermentation processes.

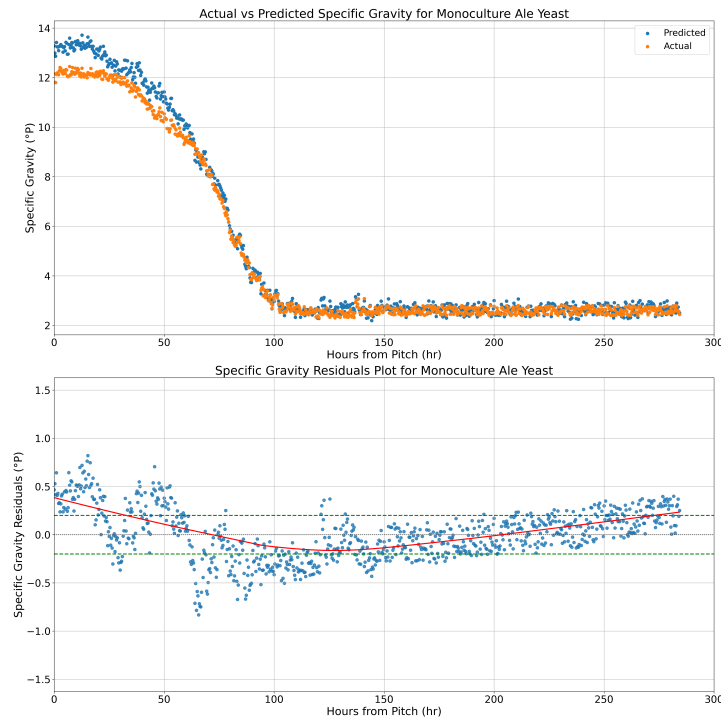


Figure 10. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowess (red traces) for specific gravity against time for the model performing against unseen ale yeast monoculture data at a primary fermentation temperature of 22 °C.

3.3.1. Monoculture Residuals

Figures 10–15 are residuals plots created from unseen fermentation data for both ale yeast (Fermentis SafAle US-05) and lager yeast (Fermentis SafLager S-23) monoculture beer fermentations, respectively; these plots were used to assess the AR-RNN’s monoculture fermentation modelling capabilities.

As may be seen from Figures 10–15, the majority of the residuals lie within the sensor tolerances of the respective attributes, indicating that the majority of the predictions are highly accurate. Furthermore, this is supported by, except for the specific gravity of the ale yeast fermentation (Figures 10–12), the lowess lines existing entirely within the instrumentation tolerances and roughly approximating zero error with some minor deviations within tolerance.

However, there are areas where the model underperforms, chiefly in temperature predictions around 120 h into the fermentation process. A likely cause of these residual spikes significantly beyond the temperature sensor’s tolerance near 120 h is the transition of the conditioning temperature from 22 °C to 5 °C after five days (120 h) of primary fermentation.

These brief spikes of error may be mitigated by feeding live temperature readings into the model or following a standard temperature profile (a recipe), as is standard practice in many breweries.

The general performance of the model for monoculture fermentations is high and favourable, with consistently low errors and few exceptions.

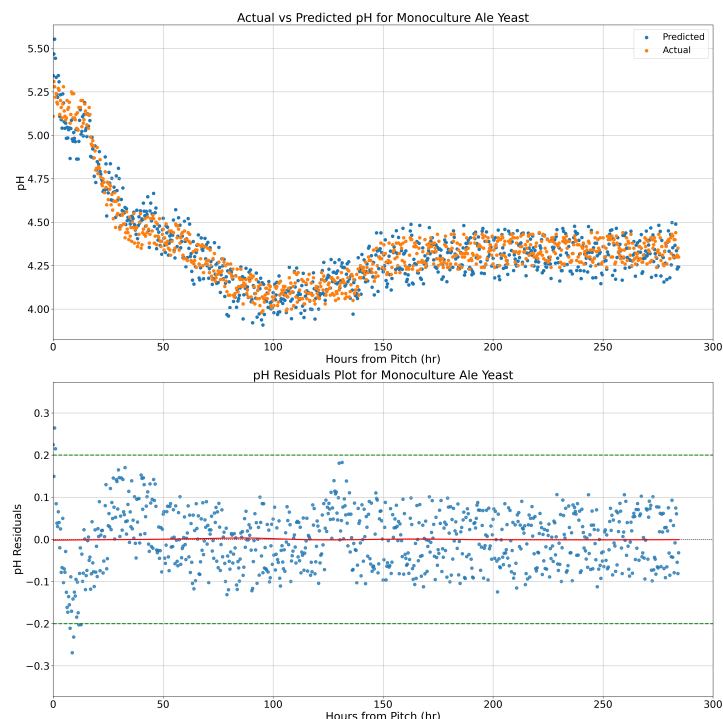


Figure 11. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowess (red traces) for pH against time for the model performing against unseen ale yeast monoculture data at a primary fermentation temperature of 22 °C.

3.3.2. Mixed-Culture Residuals

The primary focus of this research is the model's capabilities to accurately predict mixed-culture fermentation dynamics under a range of conditions. To assess the model's performance, similarly to the monoculture analysis (Section 3.3.1), residual plots were generated for the three distinct inocula mixing ratios of ale yeast (*Fermentis SafAle US-05*) and lager yeast (*SafLager S-23*) co-fermenting to produce beers with unique organoleptic profiles (Figure 4).

Figures 16–18 show the residuals plot for data from a 25:75 ratio of ale and yeast inoculum fermentation. Similarly to the temperature error spike exhibited in the monoculture fermentation model analysis (Section 3.3.1), there is also a similar spike close to the conditioning temperature changing time.

Additionally, while the lowess lines of all three data parameters agree reasonably well with an approximation of zero within instrument tolerances, the residuals provide deeper insights. The pH residuals appear broadly noisy and lie closely to the tolerance range of the instrument, indicating a good agreement between predictions and reality, as the lowess line would also suggest. However, within the specific gravity subplot, there is a noticeable oscillatory pattern between 0 and 50 h from yeast pitch within the specific gravity predictions. As these errors are only of note at the start of the fermentation process, this suggests that the model may require some additional fine-tuning to generalise the lag phase of the fermentation process.

Figures 19–21 show the residuals plot for data from a 50:50 ratio of ale and yeast inoculum fermentation. The model, in this instance, performs well overall in predicting temperature (excluding the spike during the conditioning temperature change at 120 h) and pH.

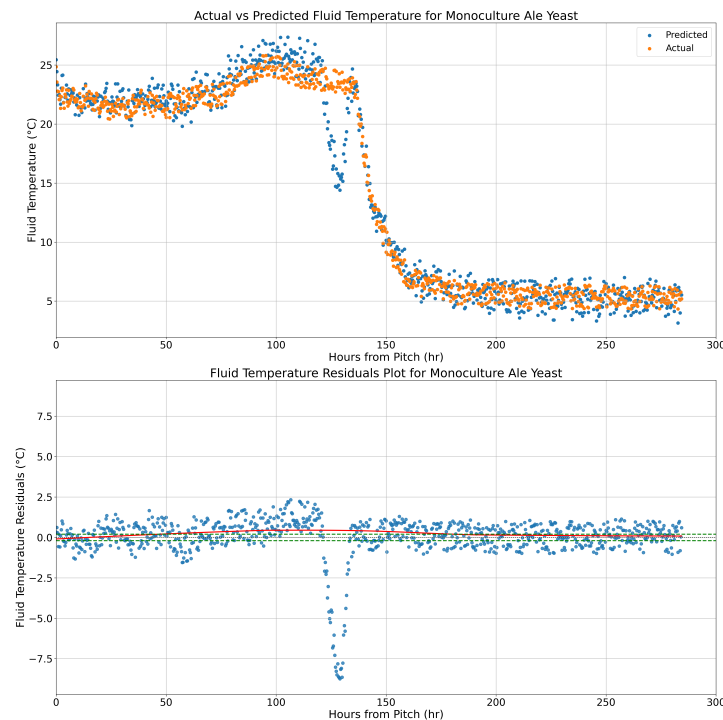


Figure 12. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowess (red traces) for fluid temperature against time for the model performing against unseen ale yeast monoculture data at a primary fermentation temperature of 22 °C.

However, the lag-phase modelling difficulties exhibited within Figures 16–18 are more pronounced in the case of Figures 19–21; while the lowess line alone does not stray too far outside of the specific gravity sensor’s tolerance, the residuals indicate that the model cannot accurately model the initial stage (0 to 50 h from yeast pitch) to a satisfactory level of reliability. Despite this, similarly to the previous instance in Figures 16–18, the model does begin to produce highly accurate predictions following the lag phase of the fermentation process.

Figures 22–24 show the residuals plot for data from a 75:25 ratio of ale and yeast inoculum fermentation. The model performs well, modelling all attributes across time with favourable lowess lines and low residuals. However, similarly to the monoculture fermentations, there is a residual spike in the temperature subplot at approximately 120 h into the fermentation process with a peak in errors approaching approximately 5 °C in magnitude.

Moreover, the specific gravity residuals are unfavourably high for a brief period between 0 and 25 h from pitch, resulting in a lowess line exceeding instrument measurement tolerances, followed by a minor over-correction by the model, leading to the lowess line approaching the limit of the lower tolerance between 100 and 125 h from pitch.

3.4. Discussion

Overall, the AR-RNN model produced could predict pH values for all five inocula combinations with a reasonably high level of fidelity, with the overwhelming majority of residuals lying within the tolerance levels of a real-world commercial pH sensor. The lowess lines approximated zero well across all pH configurations, and the residuals had few discernible patterns, indicating that the AR-RNN model had generalised the pH dynamics appropriately and could perform well as a software sensor to support monoculture and mixed-culture beer fermentation processes.

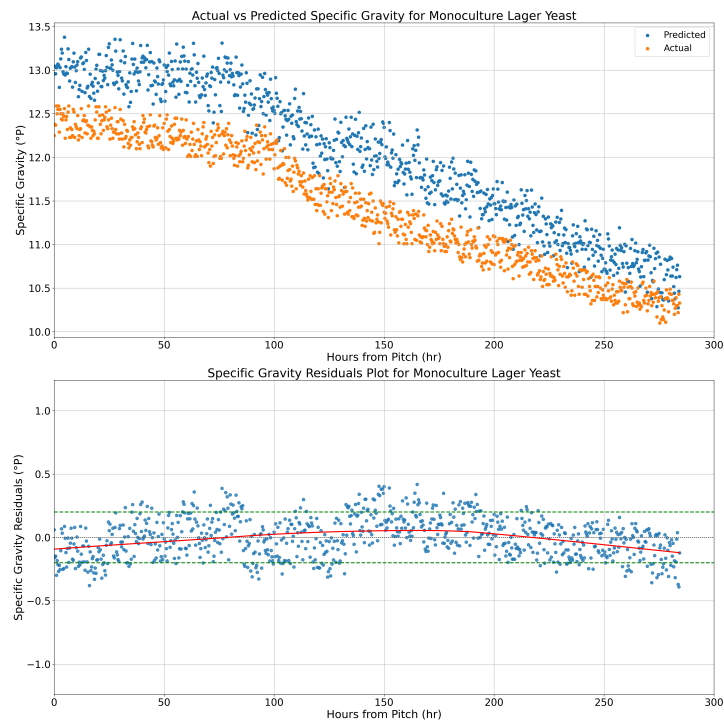


Figure 13. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for specific gravity against time for the model performing against unseen lager yeast monoculture data at a primary fermentation temperature of 22 °C.

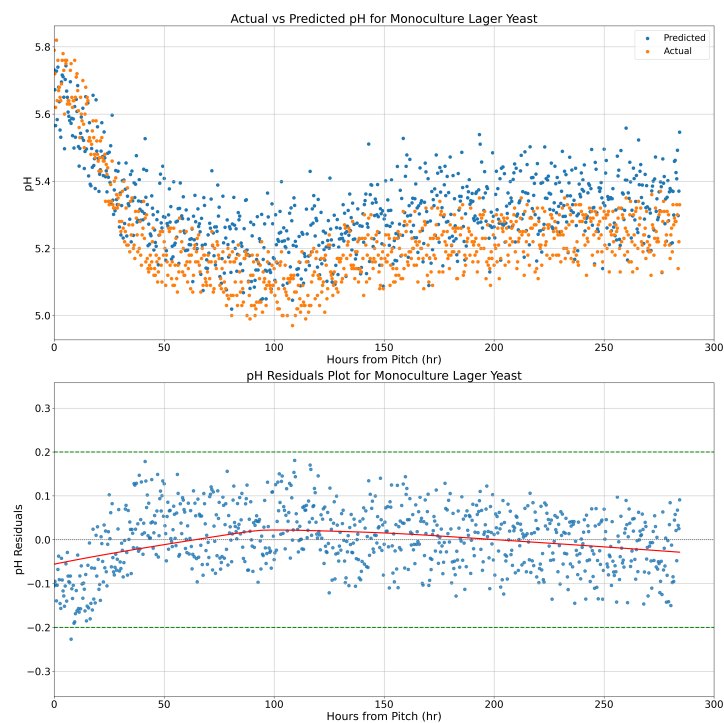


Figure 14. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for pH against time for the model performing against unseen lager yeast monoculture data at a primary fermentation temperature of 22 °C.

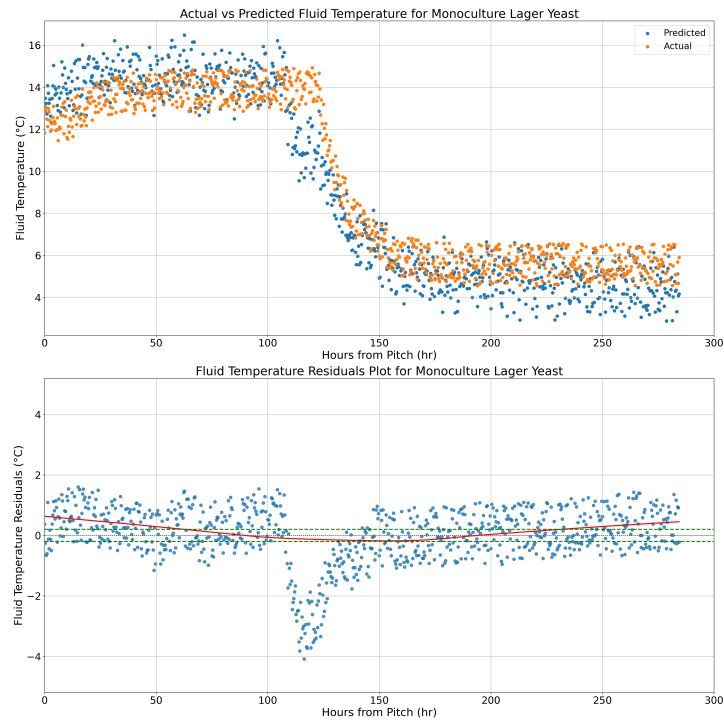


Figure 15. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for fluid temperature against time for the model performing against unseen lager yeast monoculture data at a primary fermentation temperature of 22 °C.




Figure 16. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for specific gravity against time for the model performing against unseen ale and lager yeast in a 25:75 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

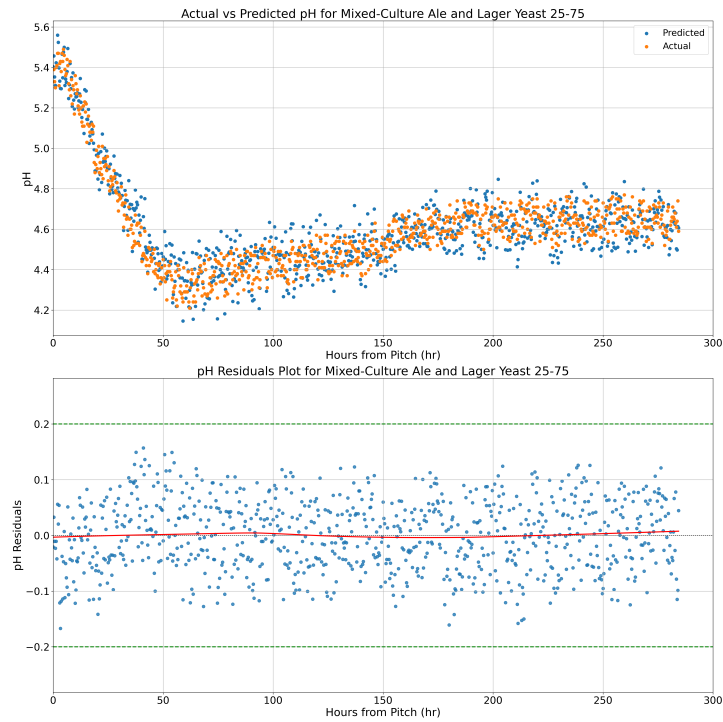


Figure 17. Actual vs. predicted values scatter plot and residuals plot (top and bottom sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for pH against time for the model performing against unseen ale and lager yeast in a 25:75 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

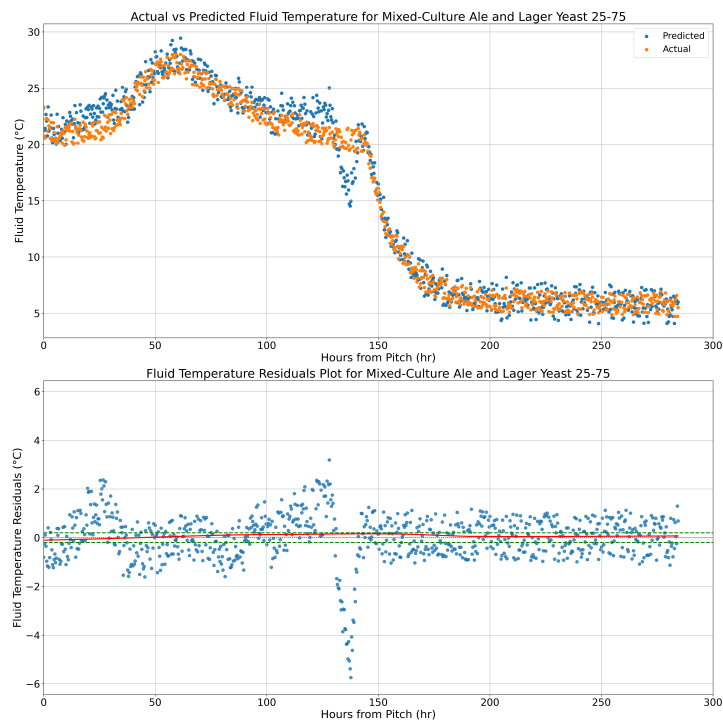


Figure 18. Actual vs. predicted values scatter plot and residuals plot (top and bottom sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for fluid temperature against time for the model performing against unseen ale and lager yeast in a 25:75 ratio mixed-culture data at a primary fermentation temperature of 22 °C.




Figure 19. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for specific gravity against time for the model performing against unseen ale and lager yeast in a 50:50 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

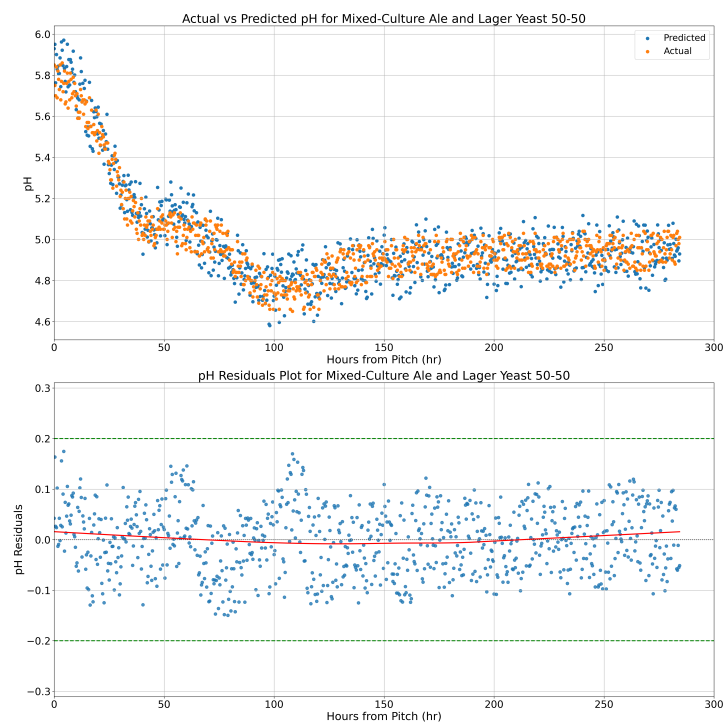


Figure 20. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for pH against time for the model performing against unseen ale and lager yeast in a 50:50 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

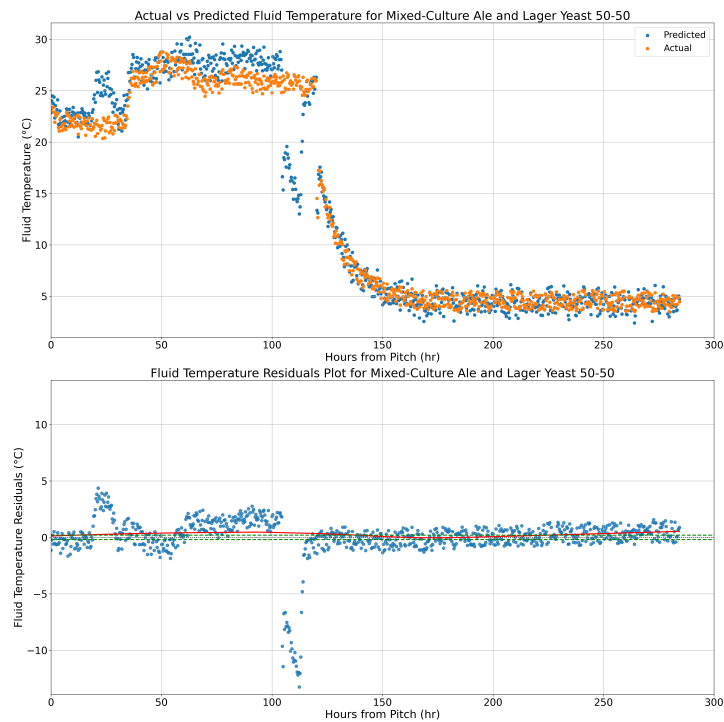


Figure 21. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for fluid temperature against time for the model performing against unseen ale and lager yeast in a 50:50 ratio mixed-culture data at a primary fermentation temperature of 22 °C.




Figure 22. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for specific gravity against time for the model performing against unseen ale and lager yeast in a 75:25 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

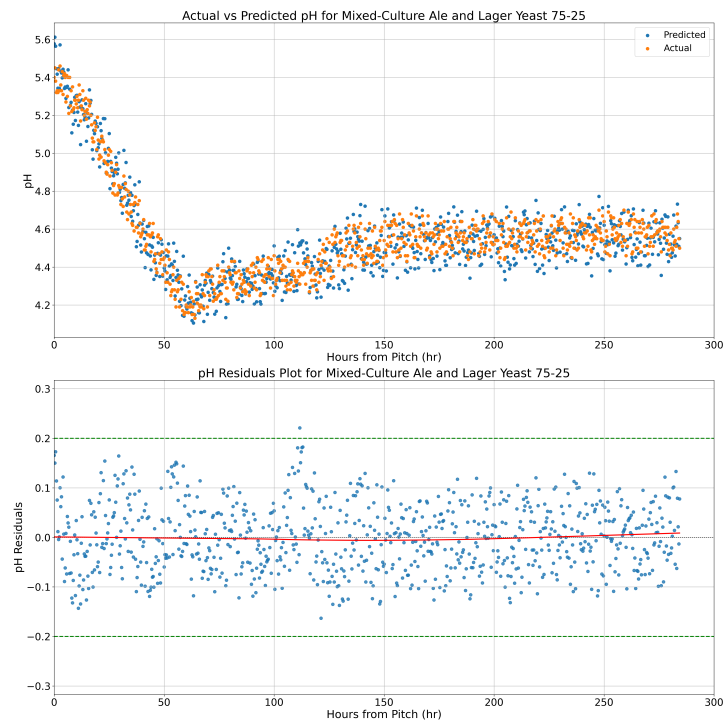


Figure 23. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for pH against time for the model performing against unseen ale and lager yeast in a 75:25 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

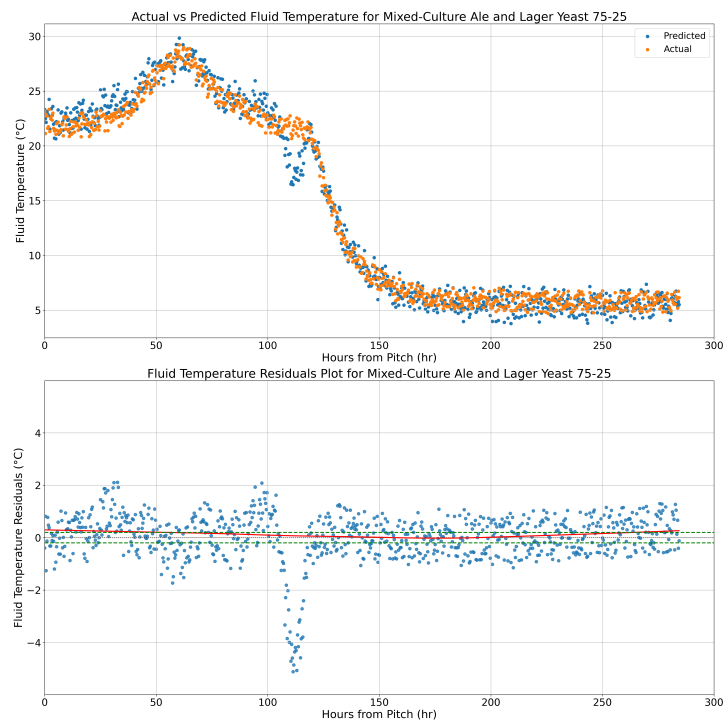


Figure 24. Actual vs. predicted values scatter plot and residuals plot (**top** and **bottom** sub-figure, respectively) with sensor tolerances (green dashes) and lines of lowest (red traces) for fluid temperature against time for the model performing against unseen ale and lager yeast in a 75:25 ratio mixed-culture data at a primary fermentation temperature of 22 °C.

The temperature analyses indicated that the fermentation temperatures were modelled reasonably accurately. However, a consistent error appeared across the residual plots at around 120 h from pitch. As mentioned before, supplementing the model's input with live temperature data or a preset profile may improve the model's overall performance. Moreover, as this spike was consistent across the testing space, it could indicate that the training data should be expanded to allow the time from yeast pitch to when the beer is cooled also to be included within the independent variable space along, with inocula and primary fermentation temperature to better-compensate for this period of the fermentation process.

The AR-RNN's predictive accuracy for specific gravity is unique among the data attributes, in that the initial predictions made by the model during the lag stage of fermentation were broadly too far outside the scope of the specific gravity sensor's tolerance to be deemed to accurately model specific gravity between 0 and 50 h from pitch. Efforts may be made to expand the training dataset further to produce a more representative sample of the lag stage of fermentation for various inocula and primary fermentation temperatures. However, as time progressed, the model did become far more accurate and may even be used to provide robust predictions of specific gravity following this initial period.

4. Conclusions

In this study, we embarked on a journey to revolutionise the modelling of mixed-culture fermentation dynamics in craft beer production. Recognising the limitations of traditional monoculture-focused models, we introduced a novel approach leveraging Autoregressive Recurrent Neural Networks (AR-RNNs). This data-driven methodology aimed to capture the intricacies of mixed-culture fermentations, offering a paradigm shift in the understanding and prediction of complex brewing processes.

Our research's key results underscore the AR-RNN model's efficacy in unravelling the mysteries of mixed-culture fermentation. We demonstrated the model's adaptability across diverse scenarios through careful data collection, preprocessing, and training. The AR-RNN successfully modelled the complexities introduced by cooperative and competitive microbial interactions, outperforming traditional models rooted in monoculture dynamics. The autoregressive aspect proved pivotal, allowing the model to forecast fermentation dynamics over extended durations accurately.

In conclusion, the AR-RNN offers a powerful tool for craft brewers seeking to innovate and push the boundaries of flavour profiles, unlocking a new era of possibilities in beer production. Furthermore, it provides a solid basis from which to delve deeper into model predictive control to improve product consistency and provide decision support to the craft brewing industry.

While this research marks a significant leap forward, there are avenues for further exploration. For instance, expanding the dataset to encompass a broader range of yeast strains, fermentation ratios, and temperatures could significantly enhance the model's accuracy and robustness. Exploring the applicability of the AR-RNN to other fermentation processes beyond beer production is another promising avenue, for instance, into the domain of biofuels. As technology and brewing practices evolve, ongoing research will be crucial to refine and expand the capabilities of data-driven models in the vibrant landscape of craft brewing. This study, we believe, is a catalyst for future studies that will continue to reshape the way we model and optimise fermentation processes in the fascinating world of craft brewing.

Author Contributions: Conceptualisation, A.O.; methodology, A.O.; software, A.O.; validation, A.O.; formal analysis, A.O.; investigation, A.O.; resources, A.O.; data curation, A.O.; writing—original draft preparation, A.O.; writing—review and editing, A.O., H.Z., and D.M.A.; visualisation, A.O.; supervision, H.Z., A.R., and D.M.A.; project administration, H.Z.; funding acquisition, H.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Informed Consent Statement: Not applicable.

Data Availability Statement: The experimental data used to support the findings of this study are available from the corresponding author upon request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ANN	Artificial Neural Network
AR-RNN	Autoregressive Recurrent Neural Network
IoT	Internet of Things
LSTM	Long Short Term Memory
RNN	Recurrent Neural Network

References

- Basso, R.F.; Alcarde, A.R.; Portugal, C.B. Could non-Saccharomyces yeasts contribute on innovative brewing fermentations? *Food Res. Int.* **2016**, *86*, 112–120. [[CrossRef](#)]
- Spitaels, F.; Wieme, A.D.; Janssens, M.; Aerts, M.; Van Landschoot, A.; De Vuyst, L.; Vandamme, P. The microbial diversity of an industrially produced lambic beer shares members of a traditionally produced one and reveals a core microbiota for lambic beer fermentation. *Food Microbiol.* **2015**, *49*, 23–32. [[CrossRef](#)] [[PubMed](#)]
- Van Oevelen, D.; Spaepen, M.; Timmermans, P.; Verachtert, H. Microbiological Aspects of Spontaneous Wort Fermentation in the Production of Lambic and Gueuze. *J. Inst. Brew.* **1977**, *83*, 356–360. [[CrossRef](#)]
- Sparrow, J. *Wild Brews: Beer Beyond the Influence of Brewer's Yeast*; Brewers Publications: Kent, OH, USA, 2005.
- Tonsmeire, M. *American Sour Beer: Innovative Techniques for Mixed Fermentations*; Brewers Publications: Kent, OH, USA, 2014.
- Shimotsu, S.; Asano, S.; Iijima, K.; Suzuki, K.; Yamagishi, H.; Aizawa, M. Investigation of beer-spoilage ability of Dekkera/Brettanomyces yeasts and development of multiplex PCR method for beer-spoilage yeasts. *J. Inst. Brew.* **2015**, *121*, 177–180. [[CrossRef](#)]
- Serra Colomer, M.; Funch, B.; Forster, J. The raise of Brettanomyces yeast species for beer production. *Curr. Opin. Biotechnol.* **2019**, *56*, 30–35. [[CrossRef](#)]
- De Roos, J.; De Vuyst, L. Microbial acidification, alcoholization, and aroma production during spontaneous lambic beer production. *J. Sci. Food Agric.* **2019**, *99*, 25–38. [[CrossRef](#)]
- Dysvik, A.; La Rosa, S.L.; De Rouck, G.; Rukke, E.O.; Westereng, B.; Wicklund, T. Microbial Dynamics in Traditional and Modern Sour Beer Production. *Appl. Environ. Microbiol.* **2020**, *86*, e00566–20. [[CrossRef](#)]
- Martens, H.; Iserentant, D.; Verachtert, H. Microbiological Aspects of a Mixed Yeast—Bacterial Fermentation in the Production of a Special Belgian Acidic Ale. *J. Inst. Brew.* **1997**, *103*, 85–91. [[CrossRef](#)]
- White, C.; Zainasheff, J. *Yeast: The Practical Guide to Beer Fermentation*; Brewers Publications: Kent, OH, USA, 2010.
- Lodolo, E.J.; Kock, J.L.; Axcell, B.C.; Brooks, M. The yeast *Saccharomyces Cerevisiae* – Main Character Beer Brewing. *FEMS Yeast Res.* **2008**, *8*, 1018–1036. [[CrossRef](#)] [[PubMed](#)]
- de Andrés-Toro, B.; Girón-Sierra, J.M.; López-Orozco, J.A.; Fernández-Conde, C. Optimization of a Batch Fermentation Process by Genetic Algorithms. *IFAC Proc. Vol.* **1997**, *30*, 179–184. [[CrossRef](#)]
- de Andrés-Toro, B.; Girón-Sierra, J.; López-Orozco, J.; Fernández-Conde, C.; Peinado, J.; García-Ochoa, F. A kinetic model for beer production under industrial operational conditions. *Math. Comput. Simul.* **1998**, *48*, 65–74. [[CrossRef](#)]
- Toh, D.W.K.; Chua, J.Y.; Liu, S.Q. Impact of simultaneous fermentation with *Saccharomyces cerevisiae* and *Torulasporea delbrueckii* on volatile and non-volatile constituents in beer. *LWT* **2018**, *91*, 26–33. [[CrossRef](#)]
- Dysvik, A.; La Rosa, S.L.; Liland, K.H.; Myhrer, K.S.; Østlie, H.M.; De Rouck, G.; Rukke, E.O.; Westereng, B.; Wicklund, T. Co-fermentation Involving *Saccharomyces cerevisiae* and *Lactobacillus* Species Tolerant to Brewing-Related Stress Factors for Controlled and Rapid Production of Sour Beer. *Front. Microbiol.* **2020**, *11*, 279. [[CrossRef](#)]
- Bhonsale, S.; Mores, W.; Van Impe, J. Dynamic Optimisation of Beer Fermentation under Parametric Uncertainty. *Fermentation* **2021**, *7*, 285. [[CrossRef](#)]
- Hepworth, N.; Brown, A.K.; Hammond, J.R.M.; Boyd, J.W.R.; Varley, J. The Use of Laboratory-Scale Fermentations as a Tool for Modelling Beer Fermentations. *Food Bioprod. Process.* **2003**, *81*, 50–56. [[CrossRef](#)]
- Peleg, M. A New Look at Models of the Combined Effect of Temperature, pH, Water Activity, or Other Factors on Microbial Growth Rate. *Food Eng. Rev.* **2022**, *14*, 31–44. [[CrossRef](#)]
- Yerolla, R.; Mehshan K M, M.; Roy, N.; Harsha, N.S.; Pavan Ganesh, M.P.; Besta, C.S. Beer fermentation modeling for optimum flavor and performance. *IFAC-PapersOnLine* **2022**, *55*, 381–386. [[CrossRef](#)]
- Paul, G.; Hawkins, C. The Real-Time Optimisation of an Industrial Fermentation Process. *IFAC Proc. Vol.* **2004**, *37*, 529–534. [[CrossRef](#)]

22. Trelea, I.C.; Titica, M.; Landaud, S.; Latrille, E.; Corrieu, G.; Cheruy, A. Predictive modelling of brewing fermentation: From knowledge-based to black-box models. *Math. Comput. Simul.* **2001**, *56*, 405–424. [[CrossRef](#)]
23. Xiao, J.; Zhou, Z.K.; Zhang, G.X. Ant colony system algorithm for the optimization of beer fermentation control. *J. Zhejiang Univ. Sci. A* **2004**, *5*, 1597–1603. [[CrossRef](#)] [[PubMed](#)]
24. Monerawela, C.; Bond, U. The hybrid genomes of *Saccharomyces pastorianus*: A current perspective. *Yeast* **2018**, *35*, 39–50. [[CrossRef](#)] [[PubMed](#)]
25. Chai, W.Y.; Teo, K.T.K.; Tan, M.K.; Tham, H.J. Fermentation Process Control and Optimization. *Chem. Eng. Technol.* **2022**, *45*, 1731–1747. [[CrossRef](#)]
26. Kondakci, T. Recent Applications of Advanced Control Techniques in Food Industry. *Food Bioprocess Technol.* **2017**, *10*, 522–542. [[CrossRef](#)]
27. Lv, N.; Bai, G.Y.; Fu, Y.J.; Yan, L.Q. Fault Diagnosis Model of Beer Fermentation Process Based on Multiway Kernel Principal Component Analysis. *Appl. Mech. Mater.* **2014**, *644–650*, 2556–2561. [[CrossRef](#)]
28. Zhu, X.; Rehman, K.U.; Wang, B.; Shahzad, M. Modern Soft-Sensing Modeling Methods for Fermentation Processes. *Sensors* **2020**, *20*, 1771. [[CrossRef](#)] [[PubMed](#)]
29. Grassi, S.; Amigo, J.M.; Lyndgaard, C.B.; Foschino, R.; Casiraghi, E. Beer fermentation: Monitoring of process parameters by FT-NIR and multivariate data analysis. *Food Chem.* **2014**, *155*, 279–286. [[CrossRef](#)]
30. Chen, L.; Nguang, S.K.; Chen, X.D.; Li, X.M. Modelling and optimization of fed-batch fermentation processes using dynamic neural networks and genetic algorithms. *Biochem. Eng. J.* **2004**, *22*, 51–61. [[CrossRef](#)]
31. Boulton, C.; Quain, D. *Brewing Yeast and Fermentation*; John Wiley & Sons: Hoboken, NJ, USA, 2008.
32. Gibson, B.R.; Lawrence, S.J.; Leclaire, J.P.R.; Powell, C.D.; Smart, K.A. Yeast responses to stresses associated with industrial brewery handling. *FEMS Microbiol. Rev.* **2007**, *31*, 535–569. [[CrossRef](#)]
33. Ginovart, M.; López, D.; Giró, A.; Silbert, M. Flocculation in brewing yeasts: A computer simulation study. *Biosystems* **2006**, *83*, 51–55. [[CrossRef](#)]
34. Gallone, B.; Steensels, J.; Prahl, T.; Soriaga, L.; Saels, V.; Herrera-Malaver, B.; Merlevede, A.; Roncoroni, M.; Voordeckers, K.; Miraglia, L.; et al. Domestication and Divergence of *Saccharomyces cerevisiae* Beer Yeasts. *Cell* **2016**, *166*, 1397–1410.e16. [[CrossRef](#)] [[PubMed](#)]
35. Bokulich, N.A.; Bamforth, C.W. The Microbiology of Malting and Brewing. *Microbiol. Mol. Biol. Rev. MMBR* **2013**, *77*, 157–172. [[CrossRef](#)] [[PubMed](#)]
36. Parapouli, M.; Vasileiadis, A.; Afendra, A.S.; Hatziloukas, E. *Saccharomyces cerevisiae* and its industrial applications. *AIMS Microbiol.* **2020**, *6*, 1–32. [[CrossRef](#)] [[PubMed](#)]
37. Baker, E.; Wang, B.; Bellora, N.; Peris, D.; Hulfachor, A.B.; Koshalek, J.A.; Adams, M.; Libkind, D.; Hittinger, C.T. The Genome Sequence of *Saccharomyces eubayanus* and the Domestication of Lager-Brewing Yeasts. *Mol. Biol. Evol.* **2015**, *32*, 2818–2831. [[CrossRef](#)]
38. Libkind, D.; Hittinger, C.T.; Valério, E.; Gonçalves, C.; Dover, J.; Johnston, M.; Gonçalves, P.; Sampaio, J.P. Microbe domestication and the identification of the wild genetic stock of lager-brewing yeast. *Proc. Natl. Acad. Sci. USA* **2011**, *108*, 14539–14544. [[CrossRef](#)] [[PubMed](#)]
39. Monerawela, C.; Bond, U. Brewing up a storm: The genomes of lager yeasts and how they evolved. *Biotechnol. Adv.* **2017**, *35*, 512–519. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.