Machine learning for process monitoring and adaptive control in continuous chromatography

Foteini Michalopoulou

October 2025

A thesis submitted to Imperial College London for the degree of Doctor of Philosophy

The Sargent Centre for Process Systems Engineering
Department of Chemical Engineering
Imperial College London, United Kingdom

Declaration of originality

I hereby certify that this is my own work and any material that is not my own work has been appropriately acknowledged and referenced. No portion of the work presented has been submitted in support of another degree or qualification at any other university or institute of learning.

Copyright statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Process efficiency and adaptability are becoming growing priorities in biopharmaceutical manufacturing. In recent years, the industry has moved towards intensified and continuous operations to meet rising demand for monoclonal antibodies and other therapeutic proteins. While such approaches promise higher productivity and reduced variability, they also introduce cyclic dynamics and non-linear behaviour that complicate downstream purification. Within this context, chromatographic separations remain a critical challenge. Mechanistic models offer detailed predictions for these systems, but their high computational cost can restrict their use in online applications. This highlights an opportunity for machine learning approaches that combine mechanistic fidelity with the computational efficiency required for real-time decision-making.

This thesis presents modelling and control frameworks that address these challenges. Firstly, the potential of data-driven models is examined, using neural networks to approximate full elution profiles and enable rapid evaluation. Building on this, a hybrid modelling approach is introduced, embedding mechanistic structure within a neural network to improve robustness while avoiding spatial discretisation. Both data-driven and hybrid models are then integrated into Bayesian optimisation workflows, demonstrating how surrogate-based methods can sustain reliable performance under online horizons where mechanistic models fail to converge. Finally, reinforcement learning is explored as a complementary paradigm for adaptive control, trained to regulate flowrate and switching times under variability, disturbances, and measurement noise.

To conclude, the presented findings establish a toolbox of machine learning methods that extend the scope of mechanistic models into online applications. The results point towards future integration of optimisation and adaptive control within digital biomanufacturing, supporting the transition to more efficient and resilient chromatographic processes.

Acknowledgements

I would like to start by expressing my deepest gratitude to my supervisor, Dr. Maria Papathanasiou, for believing in me and offering me this amazing opportunity. Her guidance, faith, and support throughout my PhD journey pushed me to believe in myself and helped me realise my full potential.

I would also like to thank Dr. Emmanouil Papadakis for his guidance and insights into the industrial world of chromatography, and Prof. Massimo Morbidelli and Prof. Ronny Pini for the time dedicated to reading this thesis, the comments and knowledge shared, and the wonderful discussion during my defence.

I feel fortunate to have been part of an amazing group, without which my PhD journey would have been very different. I would therefore like to thank my lab group, Steven, Miriam, Niki, Daniele, Alexandros, Andrea, Elena, George, Kukuh, and Alex, for the feedback and ideas, the group lunches and dinners, the conferences together, and the countless coffee breaks.

Beyond Imperial, I would like to thank my friends, old and new, for the company and friendships we have shared. I would like to thank Savvas for our daily conversations on my way to and from work, Anna and Eleftheria for our trips all over the world, and Vince for always keeping in touch despite the distance. I am also very grateful for the friends I made here, Niki, Haiting, Steven, Kostis, and Thanasis, for the board games, the discussions, and the encouragement in moments of stress and self-doubt.

I would also like to take this opportunity to thank my cousin, Katerina, who has always been more than a sister to me, for always having my back and for being my first friend, my first role model, and my sister for life.

At the heart of this journey, I would like to thank my partner, Angeliki, for loving me and supporting me, day after day, in more ways than I can name, for standing by me through the hardest moments and celebrating every success, and for taking the leap to move abroad with me to chase our dreams together.

Last but not least, I would like to thank my parents for their unwavering love and support. I would like to thank my mum, Athina, for being there for me every step of the way, always guiding me and helping me find my way, and my dad, Vangelis, for always pushing me to pursue my dreams and supporting me in all my decisions. None of this would have been possible without them.

Contents

List of Symbols **iii**

List of Figures **viii**

List of Tables **x**

1 Introduction **1**

1.1 Biopharmaceutical landscape 2

1.2 Aims & objectives 4

2 Literature review **6**

2.1 Introduction 7

2.2 Process modelling in chromatography 8

2.2.1 Mechanistic modelling 8

2.2.2 Data-driven modelling 11

2.2.3 Hybrid modelling 14

2.2.4 Remarks 17

2.3 Optimisation & control of chromatography 19

2.3.1 Model-based optimisation 19

2.3.2 Advanced process control 23

2.3.3 Remarks 27

2.4 Contributions of this work 29

3 Case study **32**

3.1 Introduction 33

3.2 The twin-column MCSGP process 34

3.2.1 The process 35

3.2.2 Mathematical modelling 36

3.2.3 System analysis 40
3.3 Remarks 44

4 Data-driven modelling **45**
4.1 Introduction 46
4.2 Methodology 48
4.2.1 Input-output structure 48
4.2.2 Synthetic data generation 50
4.2.3 Model architecture 52
4.2.4 Hyperparameter tuning 53
4.3 Results & discussion 55
4.3.1 The multiple-input/multiple-output approach 55
4.3.2 The multiple-input/single-output approach 64
4.4 Conclusions 70

5 Hybrid modelling **72**
5.1 Introduction 73
5.2 Methodology 75
5.2.1 System analysis 75
5.2.2 Model architecture 76
5.2.3 Synthetic data generation 79
5.2.4 Surrogate architecture and training 82
5.3 Results & discussion 84
5.3.1 Hybrid model validation 84
5.3.2 Interpolation 87
5.3.3 Extrapolation 90
5.4 Conclusions 92

6 Process optimisation **93**
6.1 Introduction 94
6.2 Methodology 96
6.2.1 Problem formulation 96
6.2.2 Optimisation framework 98
6.2.3 Bayesian optimisation 99
6.3 Results & discussion 101
6.3.1 Offline optimisation 101

6.3.2 Online optimisation 107
6.4 Conclusions 112

7 Reinforcement learning for adaptive control 113
7.1 Introduction 114
7.2 Methodology 116
7.2.1 Problem formulation 116
7.2.2 Control framework 118
7.2.3 Input-output structure 119
7.2.4 Reward function 120
7.2.5 Controller training 123
7.3 Results & discussion 125
7.3.1 Controller performance 125
7.3.2 Performance under feed variability 131
7.3.3 Performance against disturbances 133
7.3.4 Measurement noise 137
7.4 Conclusions 140

8 Conclusions & outlook 141
8.1 Summary of contributions 142
8.2 Main conclusions 144
8.2.1 Data-driven modelling 144
8.2.2 Hybrid modelling 145
8.2.3 Process optimisation 145
8.2.4 Reinforcement learning for adaptive control 145
8.3 Future directions 145
8.3.1 Transfer learning 146
8.3.2 Receding-horizon optimisation 146
8.3.3 Dual reinforcement learning control 147
8.3.4 Other applications 147

Dissemination of outputs 148

A Neural network architectures 171
A.1 Multiple-input/multiple-output networks 172
A.2 Multiple-input/single-output networks 173

List of Symbols

Acronyms

AEX	Anion Exchange Chromatography
ANN	Artificial Neural Network
BO	Bayesian Optimisation
CEX	Cation Exchange Chromatography
CNN	Convolutional Neural Network
CQA	Critical Quality Attribute
CSS	Cyclic Steady State
EDM	Equilibrium Dispersive Model
EI	Expected Improvement
FDA	Food and Drug Administration
GP	Gaussian Process
GPR	Gaussian Process Regression
GRM	General Rate Model
HIC	Hydrophobic Interaction Chromatography
HPLC	High-Performance Liquid Chromatography
KPI	Key Performance Indicator
LKM	Lumped Kinetic Model
LSTM	Long Short-Term Memory
mAb	Monoclonal Antibody
MAPE	Mean Absolute Percentage Error
MCSGP	Multicolumn Countercurrent Solvent Gradient Purification

MIMO	Multiple-Input/Multiple-Output
MISO	Multiple-Input/Single-Output
ML	Machine Learning
MPC	Model Predictive Control
MSE	Mean Squared Error
NLP	Non-Linear Programming
PAT	Process Analytical Technology
PDAE	Partial Differential and Algebraic Equation
PINN	Physics-Informed Neural Network
PLS	Partial Least Squares
POMDP	Partially Observable Markov Decision Process
PPO	Proximal Policy Optimisation
PSA	Pressure Swing Adsorption
QbD	Quality by Design
RL	Reinforcement Learning
RNN	Recurrent Neural Network
SMB	Simulated Moving Bed
SVR	Support Vector Regression

Continuous/Integer Variables

c	Liquid-phase concentration of component i in column h, mg/mL
c^0	Initial liquid phase concentration, mg/mL
$c^{\text{corrected}}$	Corrected liquid-phase concentration, mg/mL
c^{in}	Inlet liquid phase concentration, mg/mL
c_{av}	Average liquid-phase concentration, mg/mL
c_{data}	Liquid phase concentration predicted by the surrogate models, mg/mL
c_{feed}	Feed concentration, mg/mL
c_{mech}	Liquid phase concentration calculated by the mechanistic process model, mg/mL
H^I, H^{II}	Henry constants for adsorption sites I and II, -

Pur	Product purity, –
Q	Flowrate, mL/min
q	Solid phase concentration, mg/mL
q^*	Equilibrium solid phase concentration, mg/mL
q^I, q^{II}	Saturation capacities for adsorption sites I and II, mg/mL
r	Reward, –
t	Time, min
Y	Process yield, –
z	Axial coordinate along the column, cm

Indices

h	Column
i	Component
j	Cycle
ph	Process phase
t	Timestep

Parameters

$\alpha_1 - \alpha_8$	Experimentally defined coefficients, -
α	Weight of weak impurities in controller reward function, -
β	Weight of product in controller reward function, -
γ	Weight of strong impurities in controller reward function, -
δ	Time penalty in controller reward function, -
ε	Column porosity, –
A_{col}	Column cross-sectional area, cm^2
D_{ax}	Axial dispersion coefficient, cm^2/min
H_{mod}	Henry constant of the modifier, -
k	Lumped mass-transfer coefficient, min^{-1}
L_{col}	Column length, cm
n_{comp}	Number of components in the mixture, –
n_{col}	Number of columns, –

List of Figures

2.1 Thesis outline 31

3.1 Schematic overview of a complete cycle of the twin-column multicolumn countercurrent solvent gradient purification process. . . 35
3.2 Effect of operating condition variations on elution profiles at CSS. Each case shows one cycle with a ±50% change from nominal operating conditions in feed concentration, flowrate, or initial modifier concentration. 42
3.3 Distribution of elution product concentration values at CSS across 400 process cycles. 43

4.1 Schematic representation of the input–output structure for the synthetic data generation. 51
4.2 (a) Schematic representation of the structure of the data-driven models using MIMO ANNs, (b) Schematic representation of the structure of the data-driven models using MISO ANNs. 54
4.3 Comparison of the outlet product concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 57
4.4 Comparison of the outlet weak impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 59

4.5 Comparison of the outlet strong impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 60
4.6 Comparison of the outlet product concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 65
4.7 Comparison of the outlet weak impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 67
4.8 Comparison of the outlet strong impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples. 68

5.1 Schematic overview of the structure of the hybrid model. 78
5.2 Schematic representation of the input–output structure for the synthetic data generation. 80
5.3 Schematic representation of the input–output structure of the ANN used in the hybrid model. 82
5.4 Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under nominal operating conditions. 85
5.5 Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under interpolation conditions of feed variation: (a) −10% feed and (b) +10% feed. 87

5.6 Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under interpolation conditions of flowrate variation: (a) −20% flowrate and (b) +20% flowrate. 89
5.7 Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under extrapolation conditions of feed variation: (a) −20% feed and (b) +20% feed. 90

6.1 Optimisation workflow. 99
6.2 Convergence of Bayesian optimisation for offline optimisation of flowrates. 102
6.3 Convergence of Bayesian optimisation for offline optimisation of flowrates and initial modifier concentrations. 104
6.4 Comparison of the elution profiles of the weak impurities (blue), product (red), and strong impurities (green) at CSS, as optimised by the mechanistic (solid lines), hybrid (dashed lines) and data-driven (dash-dot lines) models in the cycle-to-cycle optimisation scenario. 109
6.5 Comparison of the elution profiles of the weak impurities (blue), product (red), and strong impurities (green) at CSS, as optimised by the hybrid (dashed lines) and data-driven (dash-dot lines) models in the minute-by-minute optimisation scenario. 111

7.1 Reinforcement learning framework. 118
7.2 Reinforcement learning agent actor-critic network structure. . . . 124
7.3 (a) Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process, the column responsible for the collection of the product and the recycling of the impure fractions, from process start-up across ten process cycles under RL control. (b) Flowrate applied by the controller on the right-hand column of the twin-column MCSGP system across the same cycles. The grey dashed lines represent the end of the first half of each process cycle, while the black dashed lines represent the completion of the second half of the cycle. 126

7.4 Comparison of the outlet concentrations of the weak impurities (blue), product (red), and strong impurities (green) for one process cycle at CSS, shown for the RL-controlled process profile (solid lines) and the nominal operating profile (dashed lines). . . . 128
7.5 Comparison of (a) outlet concentration profiles of the weak impurities (blue), product (red), and strong impurities (green), and respective (b) flowrate trajectories at CSS for RL-controlled operation under 98% (solid lines) and 95% (dashed lines) purity thresholds. 130
7.6 Comparison of (a) outlet concentration profiles of the weak impurities (blue), product (red), and strong impurities (green), and respective (b) flowrate trajectories for RL-controlled operation under two feed scenarios at CSS. Favourable conditions (solid lines) correspond to a 20% increase in product concentration and 20% decrease in impurities, while challenging conditions (dashed lines) correspond to the opposite change. 132
7.7 Outlet concentration profiles of weak impurities (blue), product (red), and strong impurities (green) together with the applied flowrate trajectory (purple) for three process cycles under RL control at CSS in the presence of one-minute disturbances. The disturbance was applied at the fourth timestep of I1 (a,b), the second timestep of B1 (c,d), and the eighth timestep of I2 (e,f). The grey dashed lines represent phase switches and the black dashed lines represent the completion of one cycle. 134
7.8 Outlet concentration profiles of weak impurities (blue), product (red), and strong impurities (green) together with the applied flowrate trajectory (purple) for three process cycles under RL control at CSS in the presence of three-minute disturbances. The disturbance was applied at the fourth timestep of I1 (a,b), the second timestep of B1 (c,d), and the eighth timestep of I2 (e,f). The grey dashed lines represent phase switches and the black dashed lines represent the completion of one cycle. 135

7.9 Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process for one process cycle under RL control, in the presence of measurement noise up to 5% applied to the pooled outlet signal (black). 138

7.10 (a) Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process, the column responsible for the collection of the product and the recycling of the impure fractions, from process start-up across ten process cycles under RL control, in the presence of measurement noise of up to 5%. (b) Flowrate applied by the controller on the right-hand column of the twin-column MCSGP system across the same cycles. The grey dashed lines represent the end of the first half of each process cycle, while the black dashed lines represent the completion of the second half of the cycle. 139

List of Tables

2.1 Indicative list of modelling contributions in chromatography. . . . 18
2.2 Indicative list of optimisation and control contributions in chromatography. . . . 28

4.1 Data generation inputs. . . . 50
4.2 Hyperparameter bounds for the Bayesian optimisation. . . . 55
4.3 Performance of the investigated multiple-input/multiple-output artificial neural networks. . . . 62
4.4 Performance of the investigated multiple-input/single-output artificial neural networks. . . . 69

5.1 Data generation inputs. . . . 79
5.2 Hyperparameter bounds for the Bayesian optimisation. . . . 83
5.3 Prediction performance metrics under nominal operating conditions. 86
5.4 Average prediction performance metrics under interpolation conditions. . . . 89
5.5 Average prediction performance metrics under extrapolation conditions. . . . 91

6.1 Optimisation bounds for decision variables. . . . 98
6.2 Offline Bayesian optimisation outputs. . . . 103
6.3 Extended offline Bayesian optimisation outputs. . . . 105
6.4 Feed concentrations for Bayesian optimisation under interpolation and extrapolation conditions. . . . 105
6.5 Bayesian optimisation outputs under interpolation and extrapolation conditions. . . . 106
6.6 Cycle-to-cycle Bayesian optimisation outputs. . . . 108
6.7 Minute-by-minute Bayesian optimisation outputs. . . . 110

7.1 Controller inputs. 119
7.2 Reward function weights. 122
7.3 Comparison of performance metrics for the RL-controlled and the nominal operating profile for one process cycle at CSS. 129
7.4 Phase durations at CSS for RL-controlled operation under 98% and 95% purity thresholds. 131

A.1 Multiple-input/multiple-output artificial neural network architectures for predicting the elution concentrations of all species. . . . 172
A.2 Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the modifier. 173
A.3 Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the weak impurities. . . 174
A.4 Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the product. 175
A.5 Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the strong impurities. . 176

Chapter 1

Introduction

1.1 Biopharmaceutical landscape

The production of high-value biopharmaceuticals has grown significantly over the past decades, driven by the expanding demand for monoclonal antibodies (mAbs), vaccines, and other therapeutic proteins (Ecker et al., 2015; Walsh, 2018). As manufacturing scales to meet global needs, the pharmaceutical sector faces increasing pressure to deliver products of consistent quality while reducing development costs and accelerating time to market (Khanal and Lenhoff, 2021; Niazi, 2025). Within this context, the purification of therapeutic proteins through chromatographic separation remains one of the most critical and resource-intensive stages of downstream processing. Ensuring high product purity and recovery, while operating under strict regulatory constraints, positions chromatography as a bottleneck in both process development and large-scale manufacturing (Cramer and Holstein, 2011; Sun et al., 2023).

In response to these pressures, the biopharmaceutical industry has increasingly turned towards continuous and intensified processing (Jungbauer et al., 2024). Continuous manufacturing offers substantial benefits, including higher productivity, improved product quality, and reduced variability, but it also introduces additional layers of complexity. Chromatographic processes illustrate this trend. Advanced multicolumn configurations enable efficient recovery of closely eluting compounds, but their cyclic operation and non-linear responses pose significant challenges for process design, optimisation, and control (Müller-Späth et al., 2013; Papathanasiou et al., 2016).

Traditional development strategies, based on iterative experimentation, are no longer sustainable in this environment. Extensive laboratory campaigns are required to adjust gradients, flow rates, or switching times, consuming large amounts of time and material while providing only limited coverage of the design space (Chan et al., 2008; Degerman et al., 2006). Mechanistic modelling has therefore become central to process development. By describing adsorption, transport, and flow phenomena from first principles, mechanistic models provide rigorous predictions of spatio-temporal concentration profiles and component interactions, supporting systematic design and scale-up (Kumar and Lenhoff, 2020; Shekhawat and Rathore, 2019). They are widely used to study column dynamics, design separation strategies, and establish regulatory confidence (Chen et al., 2020; Qu et al., 2024).

The predictive value of mechanistic models, however, can be accompanied by considerable computational demands. They are typically formulated as coupled systems of partial differential and algebraic equations (PDAEs) that require spatial discretisation and iterative solution (Kumar et al., 2021; Lin et al., 2021). In multicolumn or cyclic systems, repeated simulations are often necessary to reach steady operation, further increasing computational cost (Aumann and Morbidelli, 2007). While such models have proven invaluable for offline studies, their direct application in scenarios demanding rapid decision-making can be challenging. Online optimisation involves repeated evaluations on the timescale of process operation, and closed-loop control requires near-instantaneous responses to variability and disturbances. In both cases, the evaluation time of mechanistic models exceeds the limits imposed by industrial practice (Engell and Toumi, 2004).

This computational barrier highlights an ongoing gap between what can be modelled in principle and what can be implemented in practice. This challenge has become increasingly evident with the rise of digitalisation across the pharmaceutical sector. The Process Analytical Technology (PAT) framework, introduced through regulatory guidance by the U.S. Food and Drug Administration (FDA) (2004), has encouraged the use of advanced sensors, analytical tools, and data-driven methods to monitor and control processes in real time (Maruthamuthu et al., 2020; Shastri et al., 2023). More recently, digital twins have been proposed as virtual representations of manufacturing processes that couple models with real-time data, enabling predictive monitoring, optimisation, and scenario evaluation (Chen et al., 2020; Fuller et al., 2019; Khanal and Lenhoff, 2021). Both the PAT framework and digital twins emphasise the importance of developing models that balance predictive accuracy and computational efficiency to support real-time decision-making in chromatographic applications.

Recent progress in data science has created new opportunities to address this need. Data-driven models can be trained on synthetic or experimental datasets to approximate chromatographic dynamics at a fraction of the computational cost of mechanistic models (Boukouvala et al., 2011; Wang et al., 2017a). Once trained, they can evaluate process responses rapidly, making them suitable for online optimisation tasks. However, their reliance on training data limits their robustness when applied beyond the conditions for which they were developed (Enmark et al., 2022). Hybrid models address this limitation by embedding mechanistic

structure into data-driven frameworks, improving generalisation while maintaining computational efficiency (Ding et al., 2023; Narayanan et al., 2021a).

In parallel, reinforcement learning (RL) has emerged as a complementary paradigm for process control. By interacting with process simulations, RL agents can learn adaptive policies that optimise performance under variability and uncertainty (Kensert et al., 2024; Niezen et al., 2025). This approach is particularly relevant for chromatography, where disturbances, feed fluctuations, and measurement noise complicate conventional optimisation-based strategies. RL offers a promising route to robust, adaptive control policies capable of maintaining product quality in real time (Andersson et al., 2023; Buşoniu et al., 2018).

Taken together, these developments reflect the evolution of modelling and control in biopharmaceutical manufacturing. Mechanistic models continue to provide the foundation for process understanding, yet their computational demands limit their direct use in real-time applications. This has stimulated interest in alternative approaches, such as data-driven models that deliver rapid predictions, hybrid formulations that combine efficiency with robustness, and RL methods that learn adaptive control strategies directly from process interaction. Each of these contributes to the broader ambition of embedding reliable decision-support tools into bioprocessing workflows. Rather than replacing one another, they can be viewed as complementary, forming different parts of a toolbox that addresses the balance between fidelity, speed, and adaptability required for online operation.

It is within this context that the present thesis is positioned. The following section outlines its aims and objectives, which are directed towards the development of computational frameworks that enable online optimisation and adaptive control in continuous chromatography, bridging the gap between the descriptive accuracy of mechanistic models and the practical requirements of real-time decision-making.

1.2 Aims & objectives

The aim of this thesis is to develop machine learning-based modelling, optimisation, and control frameworks that enable online applications in continuous chromatographic separation processes. The work focuses on the twin-column multicolumn countercurrent solvent gradient purification (MCSGP) system, a

process whose cyclic operation and non-linear dynamics pose significant challenges for conventional optimisation and control, introduced in detail in Chapter 3. The central motivation is to address the computational limitations of mechanistic models, which restrict the feasibility of online optimisation and control frameworks, and to design surrogate and adaptive methods that retain predictive fidelity while meeting the time constraints required for online decision-making.

To achieve this aim, the thesis adopts a two-step approach. The first step concerns the development of surrogate models that approximate mechanistic behaviour with orders-of-magnitude reductions in computational cost. Both purely data-driven architectures and hybrid formulations that embed mechanistic knowledge are explored, with the goal of providing models that are accurate across operating regimes while suitable for rapid evaluation. The second step concerns the integration of these surrogates into optimisation frameworks, extending their role from offline simulation to real-time decision-support, alongside the development of adaptive control methods. Particular attention is given to the use of RL to design controllers capable of adapting to process variability and disturbances under stringent purity requirements.

The objectives of the thesis are pursued sequentially across the main chapters. Chapter 4 develops and benchmarks data-driven models, assessing their capacity to reproduce mechanistic simulations under nominal and interpolated conditions. Chapter 5 extends this work through the formulation of a hybrid model, which combines process information with neural network representations to improve generalisation and robustness. Chapter 6 embeds these surrogate models into optimisation frameworks, demonstrating their role in enabling both offline and online optimisation where mechanistic models are unable to converge due to their computational demands. Chapter 7 presents the development of RL-based control strategies, trained on mechanistic simulations to enable adaptive operation of the twin-column MCSGP process under process variability, disturbances, and measurement noise. Finally, Chapter 8 synthesises the contributions of the work, highlighting the implications of surrogate modelling and RL control for the broader digitalisation of bioprocessing.

Chapter 2

Literature review

2.1 Introduction

Chromatographic separation is a critical unit operation in the purification of high-value biopharmaceuticals, where product quality and process efficiency must be maintained under strict constraints (Armstrong et al., 2021). In recent years, the shift toward continuous and intensified processes has introduced additional layers of operational complexity, requiring more advanced computational tools, capable of capturing non-linear system dynamics and supporting the design, optimisation, and control of such systems. As digitalisation expands across chemical and bioprocessing industries, modelling strategies that are both accurate and computationally efficient are becoming essential for enabling model-based decision-making (Pantelides and Pereira, 2024).

First-principles models remain the dominant approach for describing chromatographic systems, offering detailed representations of mass transfer, thermodynamic interactions, and flow dynamics (Kumar et al., 2021; Shekhawat and Rathore, 2019). However, such models are often formulated as large systems of partial differential and algebraic equations (PDAEs), which are computationally expensive to simulate, particularly when involving multicolumn setups, cyclic steady state (CSS) operation, or real-time process integration (Lin et al., 2021). As a result, their direct use in real-time applications, such as process optimisation and closed-loop control, is often challenging.

In response to these challenges, recent research has focused on the development of alternative modelling strategies that retain predictive capability while reducing computational demands (Enmark et al., 2022). This thesis is motivated by the need for models that are robust to process variability, capable of generalising beyond specific operating points, and suitable for integration with optimisation and control frameworks. In this context, particular attention has been given to approaches that reduce model complexity while preserving key dynamic behaviour (Pham et al., 2023), as well as frameworks that incorporate domain knowledge into these simplified representations (Narayanan et al., 2023).

This chapter reviews the relevant literature on the modelling, optimisation and control of chromatographic separation processes, with a particular focus on approaches that aim to balance model fidelity and computational efficiency. The discussion is motivated by the challenges posed by processes such as the twin-column multicolumn countercurrent solvent gradient purification (MCSGP)

system (Aumann and Morbidelli, 2007), where cyclic dynamics, non-linearities, and tight process constraints limit the applicability of conventional methods. Emphasis is placed on identifying strategies that support accurate prediction, adaptive control, and computationally scalable optimisation.

2.2 Process modelling in chromatography

Mathematical modelling plays a central role in the analysis, optimisation, and control of chromatographic processes. A wide range of modelling approaches has been developed, reflecting different priorities across application domains (Lin et al., 2021). While some models are designed to provide mechanistic insight and support fundamental understanding of separation phenomena, others are constructed with the aim of enabling fast simulation, process monitoring, or model-based control. These varying objectives have led to a diverse modelling landscape, shaped by trade-offs between complexity, accuracy, interpretability, and computational cost.

This section provides an overview of modelling strategies used in the context of chromatographic separation. The discussion begins with mechanistic models, which remain the standard in fundamental and industrial process studies, before moving on to data-driven and hybrid methods that leverage process data or embed mechanistic assumptions into lower-order representations.

2.2.1 Mechanistic modelling

Mechanistic modelling has long served as the basis for representing chromatographic separation processes, particularly in systems where detailed understanding of transport and adsorption phenomena is required (Kumar and Lenhoff, 2020). These models are based on conservation laws and thermodynamic principles, and are typically formulated as systems of PDAEs that describe the evolution of solute concentrations across time and space. Despite significant advances in modelling frameworks and computational capabilities, mechanistic models continue to pose notable challenges in terms of numerical tractability, calibration effort, and generalisability across process conditions. Nevertheless, their role in process development, digital twins, and regulatory frameworks remains indispensable (Chen et al., 2020; Shekhawat and Rathore, 2019).

The equilibrium dispersive model (EDM) approximates chromatographic behaviour by assuming instantaneous equilibrium between the mobile and stationary phases, with axial dispersion used to capture spreading effects (Close et al., 2014; Guiochon et al., 2006; Katsoulas et al., 2023; Rehman et al., 2022). This model provides analytical tractability and is commonly used in early-stage process studies or for conceptual comparisons (Dimartino et al., 2011). However, it may lead to inaccurate predictions in systems with sharp gradients or fast-changing boundary conditions, particularly in gradient elution and multicomponent separations. Comparative studies have shown that while the EDM can approximate peak positions and widths under idealised conditions, it diverges from experimental results when non-equilibrium effects dominate (Kaczmarski et al., 2001; Qu et al., 2024; Shekhawat and Rathore, 2019).

To account for such dynamics more accurately, lumped kinetic models (LKM) introduce a linear driving force expression that relates the rate of solute uptake by the stationary phase to the deviation from local equilibrium (Felinger and Guiochon, 2004; Kluters et al., 2016). This framework improves upon the EDM by capturing some delay in adsorption dynamics while maintaining low computational cost (Felinger et al., 2004). LKM models have been successfully used in scale-up studies and optimisation tasks where full spatial resolution of the solid phase is not essential (Rehman and Qamar, 2022). However, the simplification of internal mass transfer can limit their predictive power in systems with complex adsorption kinetics or significant pore diffusion effects (Kaczmarski and Guiochon, 2007).

The most comprehensive approach is the general rate model (GRM), which resolves all major transport resistances, including film and pore diffusion, axial dispersion, and time-dependent non-linear adsorption (Shekhawat and Rathore, 2019). This model structure allows for the simulation of highly dynamic systems, such as those operating under salt gradients, pH changes, or variable loading profiles (Hahn et al., 2023). Guiochon et al. (2006) provide detailed derivations and validations of the GRM, demonstrating its ability to capture complex elution behaviours and its utility in parameter estimation workflows.

Several recent studies have advanced the use of GRMs in preparative chromatography, particularly within biopharmaceutical purification contexts (Grom et al., 2018; Hahn et al., 2016; Pia tkowski et al., 2003). Kumar et al. (2021) developed a mechanistic model for the ion-exchange purification of monoclonal

antibodies (mAbs) that incorporates salt-dependent Langmuir isotherms and non-linear mass transfer kinetics. Their work demonstrated the model's predictive accuracy under various buffer and loading conditions, but also highlighted the labour-intensive calibration process required to maintain that accuracy, especially when extrapolating to new operating points. Similarly, Pfister et al. (2017) extended GRM frameworks to oligonucleotide ion-exchange chromatography, accounting for buffer composition and pH effects explicitly. These adaptations allowed for close alignment between simulated and experimental profiles, but also introduced additional complexity in terms of parameter identifiability and solver robustness.

While the GRM offers the most physically detailed representation of chromatographic systems, this expressiveness comes at a significant computational cost. Simulation of a single column typically requires spatial discretisation using finite difference, orthogonal collocation, or finite volume methods, often resulting in hundreds of state variables (Brhane et al., 2019; von Lieres and Andersson, 2010). The problem size scales rapidly in multicolumn or cyclic systems, where the need to simulate multiple switches and recycle loops, especially until CSS is reached, can lead to runtimes on the order of hours (Müller-Späth et al., 2011). Moreover, the presence of sharp concentration fronts and strongly non-linear isotherms often induces numerical stiffness, resulting in solver convergence issues or instability unless small time steps and tight tolerances are applied (Hørsholt et al., 2019).

Beyond the computational burden of simulating high-fidelity mechanistic models, parameter estimation remains a critical bottleneck for their development and validation. Accurate specification of isotherm parameters, mass transfer coefficients, and porosity distributions requires extensive experimental data, often including overloaded breakthrough curves, multi-point gradients, and repeated calibration under varied buffer conditions (Dai et al., 2009; Yang et al., 2024). In parallel, the choice of isotherm model plays a central role in mechanistic model development (Katsoulas et al., 2025). While simplified forms such as Langmuir are frequently used for their ease of implementation, they may fail to capture the non-ideal, multicomponent, or mixed-mode interactions (Wang et al., 2017a). More complex models can offer improved fidelity but introduce additional parameters, increasing the risk of structural non-identifiability and limiting the model's reliability outside the calibration range.

In terms of model applicability, mechanistic models play a pivotal role in model-based design, process validation, and digital twin development (Silva et al., 2024; Tiwari et al., 2023; Vetter-Joss et al., 2024). Their physical interpretability and structural transparency make them attractive in regulatory settings and within Quality-by-Design (QbD) frameworks (Andris and Hubbuch, 2020; Joshi et al., 2017). They are also central to training and benchmarking data-driven and hybrid models, which often use model-generated data as synthetic inputs. Software packages such as CADET (Leweke and von Lieres, 2018) and ChromX (Hahn et al., 2015) have emerged to support scalable, modular implementation of mechanistic models, further expanding their accessibility and integration into control and optimisation workflows.

Building on these applications, mechanistic models remain the most physically detailed and interpretable approach to chromatography modelling. Their strengths lie in their ability to reveal underlying process dynamics, support predictive design, and comply with regulatory standards. However, their practical use in large-scale optimisation, real-time control, and adaptive learning remains constrained by simulation cost and calibration complexity. These limitations have driven increasing interest in alternative modelling paradigms, particularly data-driven and hybrid approaches, which are examined in the following sections.

2.2.2 Data-driven modelling

To reduce the computational burden and structural complexity of mechanistic models, purely data-driven or black-box approaches have emerged as an alternative for modelling chromatographic separation processes (Lin et al., 2021; Yang et al., 2023). These models aim to approximate system outputs directly from input–output datasets, bypassing the need for explicit knowledge of mass transfer, adsorption, or transport phenomena. Their key advantage lies in their ability to learn complex, non-linear input–output relationships from data, making them well-suited for process scenarios where experimental throughput is high, mechanistic understanding is incomplete, or model deployment speed is critical (Solle et al., 2017).

Among the most widely used data-driven approaches in chromatography are artificial neural networks (ANNs), which have been employed to predict retention times, elution profiles, and adsorption parameters across a range of chromato-

graphic modes (Anupa et al., 2025; Lin et al., 2021; Mouellef et al., 2023; Oh et al., 2007). In early work, Wang et al. (2003) applied ANNs for simulating a simulated moving bed (SMB) process, demonstrating that even relatively shallow networks could capture non-linear relationships between process inputs and key performance indicators (KPIs). Their work established a foundation for using black-box models in multicolumn systems, independent of underlying transport dynamics. Subsequent studies extended this concept to single-column systems, with Webb et al. (2009) using ANNs to predict retention behaviour under gradient elution conditions. These early efforts highlighted the capacity of ANNs to model process behaviour with minimal structural assumptions, particularly when trained on well-curated datasets.

As data availability and computational resources have expanded, more sophisticated neural network architectures have been introduced. Wang et al. (2017a) developed a multi-layer ANN framework to estimate isotherm and mass transfer parameters in protein chromatography, trained on experimental elution data collected under varying operating conditions. Their model achieved accurate profile reconstruction even in overloaded regimes, providing a data-driven alternative to inverse parameter estimation. In a subsequent study, Wang et al. (2017b) used similar architectures for fault diagnosis in chromatographic systems, identifying shifts in system behaviour, which were attributed to resin ageing and buffer variability. A related study by Liu and Papageorgiou (2019) applied data-driven modelling to guide optimal purification strategies for mAbs. The authors used neural networks to link process settings, such as pH, salt concentration, and flowrates, with product quality attributes, demonstrating the feasibility of using surrogate models to inform early-stage design decisions. These contributions illustrate how data-driven models can move beyond mere prediction, offering insight into system variability and assisting in process troubleshooting.

To capture temporal or sequential dependencies, which are particularly relevant in gradient elution or cyclic systems, recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures have been applied to time-series elution data. For instance, Li and Wang (2024) applied an LSTM network to predict full elution curves of an organic compound, achieving strong alignment with experimental data without solving mechanistic equations. Similarly, Vaškevičius et al. (2021) demonstrated that LSTM models trained on peptide sequences could outperform traditional feedforward ANNs in reversed-phase sys-

tems by capturing local dependencies in molecular structure. More recently, Mazraedoost et al. (2025) used a Transformer-LSTM model within a multitask learning framework to predict retention times of small molecules across chromatography systems, reducing the need for manual feature engineering while maintaining strong generalisability.

Parallel studies have explored convolutional neural networks (CNNs) in chromatographic modelling, primarily for tasks involving feature extraction and signal processing (Kensert et al., 2021). Matyushin et al. (2019) presented one of the earlier uses of CNNs for chromatographic data classification, applying a CNN to distinguish between types of chemical substances based on their chromatographic retention indices. More recently, Kensert et al. (2022) developed a CNN-based approach for peak detection in reversed-phase chromatograms, demonstrating improved performance under noisy conditions and overlapping peaks.

While neural network models dominate the chromatographic modelling literature, several other machine learning methods have also been explored. For instance, Gaussian process regression (GPR), with its inherent uncertainty quantification capabilities, has been used to construct response surface models for multi-step ion exchange chromatography, enabling rapid feasibility analysis and informing experimental design (Ding and Ierapetritou, 2024). Similarly, support vector regression (SVR) has been used in retention time prediction tasks, particularly in small molecule chromatography, where it showed competitive accuracy with simpler architectures (Enmark et al., 2022; Kohlbacher et al., 2006). Decision tree-based methods such as random forests and gradient boosting machines have also been applied to large retention datasets, offering robust performance with relatively low preprocessing requirements (Feng et al., 2021).

Across all methods, the successful development of data-driven models in chromatography depends on the availability of high-quality data, careful selection of input representations that capture underlying process dynamics, and robust training strategies that ensure generalisability across process conditions (Liu and Papageorgiou, 2019). Within this landscape, feedforward ANNs remain the most commonly adopted neural architecture for chromatographic modelling tasks, particularly in settings where process dynamics are encoded through physically meaningful inputs rather than inferred directly from long temporal sequences. Compared to more complex neural network structures, such as RNNs, LSTMs, and CNNs, feedforward ANNs typically involve lower training overhead and re-

duced data requirements, and they tend to exhibit more predictable convergence behaviour, which has supported their use in iterative workflows. As a result, the literature frequently favours ANN-based formulations in applications where datasets are limited, input variables are already engineered to encode process dynamics, and computational efficiency is a priority. Nevertheless, while their flexibility and low computational cost have made such models useful in development and screening tasks, they lack interpretability and extrapolation capabilities (Yang et al., 2023). These challenges have prompted the exploration of hybrid modelling approaches, which aim to integrate physical knowledge into these data-driven frameworks.

2.2.3 Hybrid modelling

Hybrid or grey-box modelling has emerged as a promising strategy for representing chromatographic processes where neither purely mechanistic nor purely data-driven approaches alone can provide the necessary balance between fidelity and computational efficiency (Agharafeie et al., 2022; McBride et al., 2020). These models aim to combine the interpretability and physical consistency of mechanistic models with the adaptability and scalability of machine learning. In the context of chromatography, hybrid formulations have been applied to accelerate simulation, facilitate parameter estimation, and improve generalisation across operating regimes (Narayanan et al., 2023; Pirrung et al., 2017). Recent efforts have focused not only on the architectural structure of the models, but also on the degree of mechanistic integration, as this directly impacts their extrapolation capability, training demands, and applicability.

Early studies in hybrid modelling of chromatographic systems include the work by Nagrath et al. (2004), who introduced a hybrid architecture that retained mechanistic mass balance equations while replacing the isotherm term with a trained neural network. The ANN received operating conditions and system variables as input and predicted adsorption behaviour, which was passed back into the mechanistic solver. This design preserved the system dynamics while enabling more flexible representation of non-linear adsorption effects, simplifying parameter estimation and allowing the model to reproduce elution profiles with reduced calibration.

A conceptually similar formulation was presented by Creasy et al. (2015),

who proposed a hybrid model that retained a lumped kinetic description of mass transfer but replaced the isotherm with a bilinear empirical interpolation scheme based on batch adsorption data. The interpolator approximated binding equilibria as a surface across protein and salt concentrations, constructed from Langmuir–Freundlich fits at fixed salt levels. The hybrid model predicted overloaded elution profiles without relying on mechanistic adsorption equations, offering strong interpolation accuracy within the experimental design space, but limited generalisability beyond it.

A comprehensive classification of hybrid models for chromatographic systems was introduced by Narayanan et al. (2021a), who proposed a structured framework for hybrid modelling in chromatography by defining models based on their degree of hybridization, from purely data-driven approaches to fully mechanistic ones. The study evaluated a data-driven model, a mechanistic model, and three hybrid neural network-based models with varying degrees of incorporated process knowledge. All models were trained to predict outlet breakthrough profiles for a protein A capture step. The results showed that even partial integration of mechanistic knowledge significantly improved prediction accuracy and generalization, especially in data-scarce or extrapolative scenarios.

Similarly, Ding et al. (2023) constructed a hybrid model for hydrophobic interaction chromatography by retaining mechanistic balance equations and enhancing the adsorption term with a neural network trained on both synthetic and experimental data. The ANN utilised the salt concentration and gradient slope as inputs and predicted adsorption coefficients, which were then used in the mechanistic PDAE system to simulate outlet concentrations. This approach mitigated the complexity of deriving the separation isotherm from mechanistic knowledge alone, by employing a simplified equation enhanced through neural network predictions.

More recently, physics-informed neural networks (PINNs) have been proposed as a distinct class of hybrid models that integrate physical laws into the training process rather than the architecture itself (Tang et al., 2023). As such, instead of explicitly solving conservation equations, PINNs incorporate these equations as soft constraints in the loss function during their training (Santana et al., 2022). A notable work by Subraveti et al. (2022a) introduced PANACHE, a PINN framework for simulating binary chromatographic processes. The model learned the full spatio-temporal concentration profiles by mapping position, time, and injec-

tion parameters to solute concentrations, while enforcing the governing PDAEs as soft constraints in the loss function. Trained solely on boundary and initial data, the PINN predicted both fluid and solid-phase concentrations across the column in a single pass, bypassing numerical solvers. While the isotherm form was fixed and not inferred, PANACHE demonstrated strong accuracy and substantial computational speed-up, highlighting the potential of physics-informed surrogates in chromatography.

The hybrid models reviewed span a range of formulations in which mechanistic and data-driven components are combined at different levels of integration. In most cases, system dynamics are retained through mass balance equations, while adsorption behaviour is approximated using neural networks or empirical surrogates. This structure allows hybrid models to reduce calibration effort and computational cost, while preserving key features of the underlying separation physics.

A key conceptual distinction within the literature concerns how physical structure is enforced. In many hybrid approaches, the governing balance equations are retained explicitly and only selected process relationships, most commonly the adsorption behaviour, are replaced or augmented by data-driven components. In these formulations, physical consistency is enforced during simulation through the mechanistic model structure, while learning is focused on those elements that are difficult to characterise accurately using first principles alone. By contrast, physics-informed approaches incorporate conservation laws directly into the training process and aim to learn full spatio-temporal concentration fields without relying on conventional numerical simulation. While this can lead to significant acceleration for forward prediction, it typically places greater emphasis on the training procedure and introduces additional complexity in terms of data requirements, loss formulation, and validation.

Importantly, these two classes of hybrid models differ not only in their mathematical formulation, but also in how and when physical constraints are enforced. Approaches that preserve explicit balance equations maintain physical consistency at every simulation step, which can be advantageous when models are embedded within iterative optimisation or control loops. PINNs, by contrast, rely on the training process to encode physical behaviour, meaning that the fidelity of the resulting model depends more strongly on the quality and coverage of the training data and on the effectiveness of the imposed constraints. As a

result, the choice of hybrid structure has direct implications for computational robustness, scalability, and applicability.

These differences indicate that hybrid model design is closely linked to the intended application. Architectural choices influence not only computational complexity, but also the degree to which physical behaviour is guaranteed during deployment and how readily the resulting model can be integrated into optimisation and control frameworks.

2.2.4 Remarks

The modelling strategies presented in this chapter reflect the evolution of chromatographic process representations from high-fidelity mechanistic descriptions to flexible data-driven and hybrid frameworks. Each class of models offers distinct advantages depending on the modelling objective (Table 2.1). Mechanistic models continue to serve as the backbone for process understanding, parameter estimation, and regulatory alignment, but are increasingly constrained by numerical stiffness, high calibration demands, and computational cost. In contrast, data-driven models offer a scalable and computationally efficient alternative, especially when applied to narrow design spaces or screening studies, though their reliance on training data and limited extrapolation restrict broader use. Hybrid models have emerged to bridge this gap, by combining physical insight with machine learning flexibility, enabling more robust predictions under uncertainty.

Across all approaches, a recurring challenge lies in finding the right balance between accuracy, speed, and interpretability, particularly in complex or cyclic systems where real-time deployment or repeated evaluations are required. This has driven recent efforts toward more modular frameworks, reduced-order representations, and learning-based surrogates that retain mechanistic structure. Still, many of these models remain difficult to apply consistently, highlighting the need for methods that can adapt to different data regimes and modelling goals.

Ultimately, modelling in chromatography is not an end in itself, but a tool for decision-making. Its value depends on how well it supports tasks such as optimisation, control, and process design. The next section explores how these modelling strategies can be embedded into such frameworks, with a focus on their algorithmic structure and application to cyclic and multicolumn operations.

Table 2.1: Indicative list of modelling contributions in chromatography.

Model	Authors	Contribution	Method	Key Points
Mechanistic	Felinger and Guiochon (2004)	EDM vs LKM; kinetic effects	Linear driving-force model	Improved accuracy of LKM - simplified diffusion
	Kaczmarski and Guiochon (2007)	Detailed mass-transfer kinetics	Finite-volume simulation	Quantified diffusion - high computational cost
	Müller-Späth et al. (2011)	Model–experiment validation of MCSGP	Simulation–experiment study	Accurate yield & purity prediction
	Kumar et al. (2021)	Ion-exchange model for mAbs	Salt-dependent Langmuir isotherm	Accurate profiles - intensive calibration
	Hahn et al. (2023)	GRM for gradient systems	Dynamic adsorption & dispersion	Realistic gradients - solver stiffness
Data-driven	Wang et al. (2003)	ANN for SMB simulation	Feedforward neural network	Learns non-linear patterns - weak extrapolation
	Webb et al. (2009)	ANN for gradient elution	Input–output training	Predictive & flexible - data-dependent
	Wang et al. (2017a)	ANN for parameter estimation	Multi-layer network	Accurate profiles - large data requirements
	Liu and Papageorgiou (2019)	Surrogate for mAb purification	Models linking pH, salt, & flowrate	Fast optimisation - limited range
	Mazraedoost et al. (2025)	Transformer–LSTM for retention	Sequence-based learning	Broad generality - high training cost
Hybrid	Nagrath et al. (2004)	Hybrid SMB model with ANN	Mass balance & ANN isotherm	Retains dynamics - coupling effort
	Creasy et al. (2015)	Bilinear adsorption interpolation	Lumped kinetics & empirical isotherm	Local accuracy – weak extrapolation
	Narayanan et al. (2021a)	Hybridisation framework	Mechanistic vs ANN comparison	Improved accuracy - high training cost
	Subraveti et al. (2022a)	PANACHE framework	PDE-constrained PINN	Fast simulation - heavy training
	Ding et al. (2023)	Hybrid HIC model	ANN adsorption in PDAE system	Simplified isotherm - needs synthetic data

2.3 Optimisation & control of chromatography

Optimisation and control of chromatographic processes are essential to ensure high product quality, maximise yield, and maintain efficient resource utilisation (Engell and Toumi, 2004). These tasks are particularly demanding in the context of complex separation systems, such as gradient elution and multicolumn setups, where non-linear dynamics, cyclic behaviour, and tight operational constraints can challenge conventional design methods. As bioprocessing moves toward continuous, intensified, and digitally integrated paradigms, model-based strategies have gained prominence, enabling systematic exploration of the design space and real-time adaptation to variability.

This section surveys the literature on model-based optimisation and control strategies for chromatographic separation processes. Emphasis is placed on the algorithmic structure and performance of different approaches, with particular attention to how modelling choices influence the tractability, reliability, and scalability of the resulting frameworks.

2.3.1 Model-based optimisation

Model-based optimisation provides a systematic framework for improving chromatographic separations with respect to performance metrics such as product yield, purity, productivity, cycle time, and buffer consumption (Kawajiri, 2021). The choice of optimisation strategy is tightly linked to the underlying model used to represent the process. Mechanistic models offer physical consistency and predictive accuracy but are often computationally demanding. Data-driven and hybrid models, by contrast, reduce evaluation cost and are amenable to more advanced global search methods, albeit with trade-offs in interpretability and extrapolation. This section presents a review of model-based optimisation strategies applied to chromatographic systems, with a particular focus on how different models are leveraged to enable efficient and reliable optimisation.

2.3.1.1 Mechanistic model-based optimisation

Mechanistic models remain the foundation for chromatographic process optimisation, particularly in applications where accuracy and regulatory confidence are paramount (Kawajiri, 2021; Ng et al., 2012). One of their key advantages lies

in their ability to capture detailed process interactions and extrapolate beyond the experimental design space (Hahn et al., 2023; Kawajiri and Biegler, 2006). However, this expressiveness comes at a computational cost. The stiffness and dimensionality of PDAE-based models, particularly in multicomponent or gradient systems, can pose challenges for both simulation and optimisation (Sommerfeld and Strube, 2005). As a result, most contributions in this space have focused on offline optimisation, where computational time is not a limiting factor, and design decisions can be made in advance of operation.

Several studies have demonstrated the utility of mechanistic models in process intensification and throughput optimisation. Taguado Menza et al. (2024), for instance, proposed a model-based optimisation framework for oligonucleotide purification using step-gradient elution. Their mechanistic model was used to optimise cycle time, productivity, and buffer use under strict purity constraints. By adjusting gradient slopes and step durations, the authors demonstrated how systematic design choices could significantly enhance throughput while maintaining regulatory compliance.

Similarly, Degerman et al. (2006) applied constrained optimisation to a two-component preparative chromatography step using a GRM with Langmuir kinetics. Decision variables included cut points, loading volume, and flowrate, with the objective of maximising productivity and yield under a purity constraint. The study demonstrated that framing purity as a direct optimisation constraint, rather than a post hoc filter, improved convergence and enabled a more efficient search of the design space.

The choice of model also directly influences the quality of optimisation results. In that context, Osberghaus et al. (2012) investigated the impact of model fidelity on optimal operating conditions by comparing mechanistic and empirical formulations for a ternary protein separation. While empirical models offered faster convergence, only the mechanistic model correctly predicted the elution order and resolved overlapping peaks. The optimisation targeted flowrate profiles and elution times to maximise separation efficiency while reducing solvent consumption, highlighting how physically grounded models can improve decision quality even at the cost of added complexity.

Other contributions have focused on incorporating quality constraints directly into the optimisation problem. Teoh et al. (2001) formulated a dynamic optimisation framework for ion-exchange chromatography under strict product

specifications. Using control vector parameterisation, they optimised the loading and elution stages to minimise cycle time while satisfying predefined yield and purity thresholds. Their results illustrated how mechanistic optimisation can be used to enforce critical quality attributes (CQAs), though they also noted numerical challenges due to the non-convex nature of the solution space.

In more complex systems, mechanistic models have been applied to explore trade-offs across multiple objectives. Kawajiri and Biegler (2008) developed a non-linear programming (NLP) framework for the optimal zone configuration of SMB processes. Their formulation targeted feed throughput while satisfying purity and recovery constraints, and examined the trade-off between productivity and desorbent usage by modifying scalarised objective functions. While the study did not construct a formal Pareto front, it provided valuable insight into how economic and operational objectives interact in continuous separations. A related study by Chan et al. (2008) used mechanistic optimisation to compare single- and multi-column processes under different economic scenarios. Their work focused on maximising net profit while meeting product specifications, demonstrating how mechanistic models can inform techno-economic decisions across alternative process configurations.

More recent contributions have explored ways to reduce the cost of mechanistic optimisation by integrating data-efficient global search strategies. Silva et al. (2024) implemented a Bayesian optimisation framework to accelerate high-throughput chromatographic process development, combining automated experimentation with a mechanistically calibrated digital twin. Rather than replacing the model, Bayesian optimisation was used to guide experiment selection and operating point exploration, targeting improvements in throughput and product quality. Resin lifetime was modelled to support decision-making over long-term productivity, rather than directly optimised. This study illustrates how probabilistic global search methods can enhance mechanistic optimisation workflows, particularly when integrated into digital infrastructures that support iterative experimentation.

While mechanistic models provide unparalleled physical insight, their use in online optimisation is limited by the associated high computational expense, especially in cyclic or multicolumn systems. Such challenges have motivated the adoption of surrogate and hybrid models, which reduce evaluation time while preserving predictive fidelity within the training domain.

2.3.1.2 Surrogate model-based optimisation

Surrogate-based optimisation has emerged as a means to accelerate decision-making in chromatographic systems (Pham et al., 2023; Pirrung et al., 2017). These approaches involve the construction of computationally efficient models that approximate the input–output behaviour of high-fidelity mechanistic simulations. Once trained, the surrogate replaces the original model in the optimisation loop, enabling faster convergence and broader design space exploration (Lin et al., 2021).

One of the earliest demonstrations of surrogate-assisted optimisation in chromatography was presented by Nagrath et al. (2004), who embedded a trained ANN into a mechanistic optimisation framework. Their study optimised step times, feed concentrations, and gradient profiles to improve product yield and resolution, offering a computationally efficient alternative to purely mechanistic search strategies. Webb et al. (2009) later applied ANNs to identify optimal gradient elution conditions in high-performance liquid chromatography (HPLC). The optimisation task focused on gradient slope and solvent composition to maximise peak separation across multiple analytes. By bypassing iterative experimentation, the ANN enabled direct condition selection within a constrained design space.

At the process flowsheet level, Keulen et al. (2023) developed ANN surrogates to support optimisation of downstream purification trains involving anion exchange (AEX), cation exchange (CEX), and hydrophobic interaction chromatography (HIC) units. The decision variables included unit selection, sequence, and operating conditions, with the optimisation targeting a scalarised performance index combining yield, purity, and buffer usage. Constraints related to resin compatibility, number of steps, and minimum recovery were incorporated, and surrogate models were used to accelerate global screening prior to mechanistic validation.

In a similar context, Liu and Papageorgiou (2018) addressed multi-objective optimisation in mAb purification using ε-constraint methods. Their framework simultaneously optimised column sizing, unit ordering, and operating points to minimise cost of goods while maintaining product quality and recovery within acceptable limits. The surrogate models enabled rapid trade-off analysis across competing design objectives, supporting decision-making in early-stage process

development.

More recent work has sought to refine surrogate-based optimisation pipelines for higher fidelity and efficiency. For instance, Ding and Ierapetritou (2024) introduced a GPR-based framework for multistep ion-exchange chromatography, optimising gradient slopes, elution timing, and salt concentrations to maximise productivity under constraints on yield, purity, and salt load. The workflow employed a two-phase strategy, by first learning the feasible region via constraint surrogate modelling and then using adaptive Bayesian optimisation to identify the global optimum. The framework achieved a 50.1% improvement in productivity and outperformed conventional genetic algorithms in both convergence speed and sample efficiency.

Building on their physics-informed surrogate framework, Subraveti et al. (2022a) applied their trained PINN model to the optimisation of binary chromatography. Decision variables included injection volume and gradient slope, with the objective of maximising resolution while ensuring constraints on purity and breakthrough. By embedding a PINN into an optimisation loop, they demonstrated how physics-based surrogates could support multi-objective design under tight physical and operational constraints, without requiring explicit numerical solvers during optimisation.

Together, these contributions illustrate the growing role of surrogate models in enabling efficient and scalable optimisation across a range of chromatographic applications. Whether applied to early-stage screening, multi-objective trade-offs, or flowsheet design, surrogate-based frameworks reduce the computational bottlenecks associated with mechanistic models while retaining process relevance. Importantly, their speed and adaptability make them especially suited to dynamic environments where optimisation must occur continuously or in response to real-time disturbances. As such, surrogate models have become foundational not only for offline optimisation, but also for enabling closed-loop integration with advanced control strategies, which will be explored in the following section.

2.3.2 Advanced process control

As chromatographic processes evolve toward continuous and intensified configurations, traditional open-loop operation becomes insufficient to ensure product quality and process efficiency under variable conditions (Lin et al., 2021). Online

adaptation to disturbances such as resin ageing, feed variability, and equipment inconsistencies is particularly critical in multicolumn and cyclic systems, where dynamic behaviour and tight performance constraints complicate process operation. In response, a variety of control strategies have emerged, ranging from classical PID loops to advanced model-based and learning-driven algorithms, reflecting the growing importance of closed-loop integration in biopharmaceutical separations (Armstrong et al., 2021; Krättli et al., 2011; Küpper and Engell, 2008; Papathanasiou et al., 2019a).

Historically, PID controllers have served as the standard feedback strategy for biopharmaceutical processes, including chromatography (Armstrong et al., 2021). These controllers operate by correcting deviations from a target signal, often based on UV, pH, or conductivity measurements, and remain widespread due to their simplicity and tuning ease. Ündey et al. (2010) reviewed real-time monitoring practices in biomanufacturing and documented the application of PID loops to regulate critical parameters during purification, particularly in early-stage continuous operations. While effective for narrow operating ranges, PID control lacks the flexibility and constraint handling required in multivariable, non-linear systems.

A notable application of PID in chromatography was presented by Krättli et al. (2013b), who implemented an online controller for the MCSGP process. The system leveraged at-line HPLC measurements to adjust cut points and product collection times, compensating for disturbances in feed composition and flowrate. Although PID tuning was process-specific and non-adaptive, the method demonstrated stable control of product yield and purity under process variability.

To address these limitations, model-based strategies, particularly model predictive control (MPC), have been developed for both batch and continuous chromatographic processes. MPC offers the advantage of constraint handling and predictive horizon optimisation, making it well suited for managing process transitions and productivity trade-offs. One of the earliest model-based control strategies for chromatographic systems was introduced by Natarajan and Lee (2000), who proposed a repetitive MPC formulation for SMB chromatography. The controller leveraged the system's cyclic nature by updating operating trajectories after each cycle using simulated concentration feedback and a reduced-order model. This approach demonstrated improved robustness to model mismatch

and enhanced product yield.

This concept was further developed by Grossmann et al. (2008), who proposed an experimentally validated cycle-to-cycle control strategy for SMB, where switching times were adjusted based on outlet concentration averages. The same group later extended the framework to the MCSGP process (Grossmann et al., 2010). In their study, they used a reduced-order mechanistic model to perform offline optimisation of gradient and switching profiles, demonstrating how model-based control could exploit the degrees of freedom offered by the MCSGP system's structure.

Other contributions have since explored more sophisticated model formulations. For instance, Lee and Seidel-Morgenstern (2018) developed a non-linear MPC controller for SMB that incorporated adsorption and transport dynamics using a mechanistic state-space model. The study compared linear and non-linear formulations and demonstrated superior separation performance with the latter, albeit with increased computational demands, thus limiting its applicability in real-time settings.

While these contributions focused on centralised predictive control, other studies investigated structurally simplified or adaptive approaches to improve tractability and real-time applicability. In that context, Papathanasiou et al. (2016) proposed a decentralised control strategy for the twin-column MCSGP process, where each column was described by a separate state-space model. Control was applied independently to each unit, with modifier concentration selected as the primary manipulated variable based on sensitivity analysis. The simulation results showed improved robustness and disturbance rejection compared to traditional PID control.

Complementing this, Steinebach et al. (2016) developed an adaptive model-based control strategy for a two-column Protein A affinity capture system. Unlike classical MPC, their approach dynamically adjusted switching times and flowrates based on process state estimates derived from a mechanistic model. The strategy was validated experimentally, and it successfully improved resin utilisation, buffer efficiency, and yield compared to static control policies.

In parallel, broader efforts have aimed to incorporate predictive control into digital infrastructure for process understanding and monitoring. Tiwari et al. (2023) presented a digital twin framework for continuous mAb purification, integrating mechanistic simulations with Process Analytical Technology (PAT) data

to enable in silico testing of control strategies and support design space exploration. While not used in closed-loop control, the framework laid groundwork for real-time advisory systems. Similarly, Engell and Toumi (2004) proposed an optimisation-enhanced feedback control strategy for binary chromatography, where offline-optimised switching policies were dynamically corrected based on feed variability.

Alongside these model-based frameworks, several contributions have explored model-free, heuristic or statistics-based control strategies. Early contributions explored rule-based and empirical control schemes, such as the work of Suvarov et al. (2012), who implemented a PI controller for SMB systems derived from wave theory rather than explicit process models. This approach maintained product purity under disturbances through cycle-wise updates that required little process knowledge. Building on this concept, later studies integrated real-time analytics and machine learning to enhance adaptability. Rüdt et al. (2017) implemented a soft-sensor-based control strategy using inline UV/Vis spectroscopy and partial least squares (PLS) modelling to dynamically terminate the load phase in Protein A capture. The controller, validated on a continuous setup, demonstrated improved resin utilisation and reduced product loss, illustrating how PAT tools can support feedback control. In a similar vein, Xie et al. (2023) applied machine learning to develop a self-optimising controller for capture chromatography, using UV/NIR signals to train a predictive model that dynamically updated column settings. Fioretti et al. (2024) extended this concept to an event-driven UV-based scheme for MCSGP, wherein peak boundaries were detected using dynamic signal thresholds to trigger product and recycle decisions. These approaches offer robustness and speed in low-data or poorly modelled regimes, though they may struggle with long-term adaptability or unseen conditions.

Building on these developments, reinforcement learning (RL) has emerged as a promising approach for closed-loop control of chromatographic systems, particularly those characterised by complex switching logic, non-linearity, and operational constraints (Kensert et al., 2024; Nikita et al., 2021). Unlike MPC, RL does not require an explicit process model, instead learning a policy that maps process states to control actions through trial-and-error interaction with a simulation or environment. This makes RL particularly attractive in cyclic systems, where predefined setpoints or reference trajectories are difficult to specify. Early studies by Andersson et al. (2023) have applied RL to SMB systems, demon-

strating its ability to learn switching rules and flowrate profiles that improved productivity, albeit at the cost of lower product purity.

Together, these studies demonstrate how control strategies for chromatography have progressed from fixed, rule-based approaches to increasingly adaptive, model-driven, and data-informed frameworks. While many contributions remain at the simulation level, the integration of real-time data, reduced-order models, and learning-based policies offers a path forward for enabling closed-loop operation in complex and variable processes.

2.3.3 Remarks

The optimisation and control strategies reviewed in this chapter reflect a shift toward greater adaptability, speed, and scalability in chromatographic decision-making (Table 2.2). Mechanistic models are the foundation of high-fidelity optimisation workflows and remain indispensable in constrained or high-stakes settings. However, their computational demands and calibration requirements pose limitations for real-time use, especially in cyclic or multicolumn systems.

Surrogate-based approaches, both data-driven and hybrid, offer promising alternatives. Their ability to approximate key process dynamics at lower computational cost has enabled their use in design space exploration, multi-objective optimisation, and closed-loop applications. When properly structured, these models can maintain predictive reliability across varying operating regimes and support integration with constrained or probabilistic search methods.

In the context of control, the literature highlights a parallel evolution. Early cycle-to-cycle and MPC strategies have been expanded through the use of nonlinear formulations, decentralised architectures, and adaptive schemes. While few have progressed to full-scale deployment, several have demonstrated feasibility in experimental settings. Reinforcement learning has recently emerged as a flexible model-free alternative, well-suited to problems characterised by high non-linearity and cyclic operation.

Across both optimisation and control, the role of the models remains central. As chromatographic processes become more integrated and variable, the ability to make fast, reliable decisions under uncertainty will depend on the continued development of modelling frameworks that balance accuracy, speed, and generalisability.

Table 2.2: Indicative list of optimisation and control contributions in chromatography.

Approach	Authors	Contribution	Method	Key Points
Mechanistic optimisation	Teoh et al. (2001)	Dynamic optimisation of ion-exchange	Control vector parameterisation	Minimised cycle time - enforced CQAs
	Degerman et al. (2006)	Purity-constrained optimisation	GRM with Langmuir kinetics	Maximisation of productivity & yield
	Kawajiri and Biegler (2008)	NLP design for SMB zones	Non-linear programming	Throughput–solvent trade-offs
	Osberghaus et al. (2012)	Model fidelity vs optimal conditions	Mechanistic vs empirical comparison	Mechanistic accuracy - higher cost
	Taguado Menza et al. (2024)	Step-gradient optimisation	Mechanistic model with purity constraint	Higher throughput - lower buffer use
Surrogate optimisation	Nagrath et al. (2004)	Hybrid ANN-assisted optimisation	Mechanistic & ANN surrogate	Faster search - retained physics
	Webb et al. (2009)	ANN design of gradient elution	Feedforward ANN	Direct condition pick - fewer trials
	Liu and Papageorgiou (2018)	Multi-objective mAb purification	ε-constraint strategy	Cost vs quality under constraints
	Keulen et al. (2023)	Flowsheet optimisation	ANN surrogates & multi-objective search	Yield–purity–buffer trade-offs
	Ding and Ierapetritou (2024)	GPR & BO for ion-exchange	Gaussian processes & adaptive BO	Data-efficient - large productivity gains
Control strategies	Natarajan and Lee (2000)	Repetitive MPC for SMB	Reduced-order model with cycle updates	Improved robustness - higher yield
	Grossmann et al. (2010)	Model-based control of MCSGP	Reduced-order model & scheduling	Optimised switching - offline planning
	Krättli et al. (2013b)	PID for MCSGP with at-line updates	At-line HPLC & PID tuning	Stable yield & purity - non-adaptive
	Rüdt et al. (2017)	PAT-based soft - sensor control	Inline UV/Vis & PLS modelling	Dynamic load control - improved resin use
	Andersson et al. (2023)	Adaptive RL control of single column	Actor–critic policy trained in silico	Productivity-purity trade-off

2.4 Contributions of this work

The review presented in this chapter has shown that a broad range of modelling approaches has been developed for chromatographic separation processes, with increasing attention to computational efficiency and data integration. Within this context, mechanistic, data-driven, and hybrid models have each contributed to addressing different aspects of complexity in chromatographic system behaviour, while a growing body of work has explored their application in process optimisation and control. Across these developments, model selection, system dynamics, and integration into decision-making frameworks have been central considerations.

The first knowledge gap that emerges concerns the development of surrogate models for cyclic chromatographic processes. Most data-driven approaches to date have been constructed for single-column or batch systems and are trained to predict process parameters and KPIs, such as purity, recovery, and yield. Few contributions have addressed the challenge of capturing cyclic elution dynamics using only measurable process quantities. As such, there is a need to develop data-driven modelling strategies that are computationally efficient, can be used despite limited data availability, and can support deployment in dynamic optimisation and control tasks.

A second knowledge gap relates to the development of hybrid models suitable for real-time applications in chromatographic processes. While hybrid modelling approaches have shown promise in reducing complexity and improving extrapolation capabilities, most rely on computationally expensive discretised PDAE systems, limiting their deployability in real-time applications. Few contributions have proposed modelling architectures that preserve core mechanistic structure while enabling fast and robust predictions based solely on observable variables. As such, there is a need for hybrid formulations that combine physical insight with data-driven components to retain interpretability and efficiency across varying operating conditions.

A third gap is identified in the area of online model-based process optimisation. While mechanistic models have been widely used for process optimisation, their computational complexity and simulation time requirements often limits their applicability in real-time settings. Existing data-driven and hybrid models have shown potential in reducing simulation times, yet most formulations

are either decoupled from optimisation or limited to static, offline scenarios. As such, there is a need for surrogate models that are not only accurate and computationally efficient, but can also be seamlessly integrated within optimisation frameworks capable of operating at different temporal resolutions during process operation.

A final knowledge gap emerges in the development of adaptive control strategies for chromatographic separation processes. While model-based and rule-driven methods have been applied to multicolumn systems, existing approaches are often limited to fixed schedules or rely on reduced-order models with limited flexibility. In processes characterised by non-linear behaviour, partial observability, and time-varying disturbances, there is a need for control frameworks that can dynamically adapt to system evolution without predefined reference trajectories.

In this thesis, these knowledge gaps are addressed in the following chapters (Figure 2.1):

- Chapter 4: Data-driven modelling approaches
- Chapter 5: Hybrid model development
- Chapter 6: Model-based optimisation
- Chapter 7: Reinforcement learning for adaptive control

This thesis presents a set of digital tools designed to support data-driven and adaptive decision-making in complex continuous chromatographic systems. The methodologies presented throughout the thesis are developed using the twin-column MCSGP process as a case study, which is presented in Chapter 3. However, the proposed frameworks are designed to support broader application to continuous separation systems.

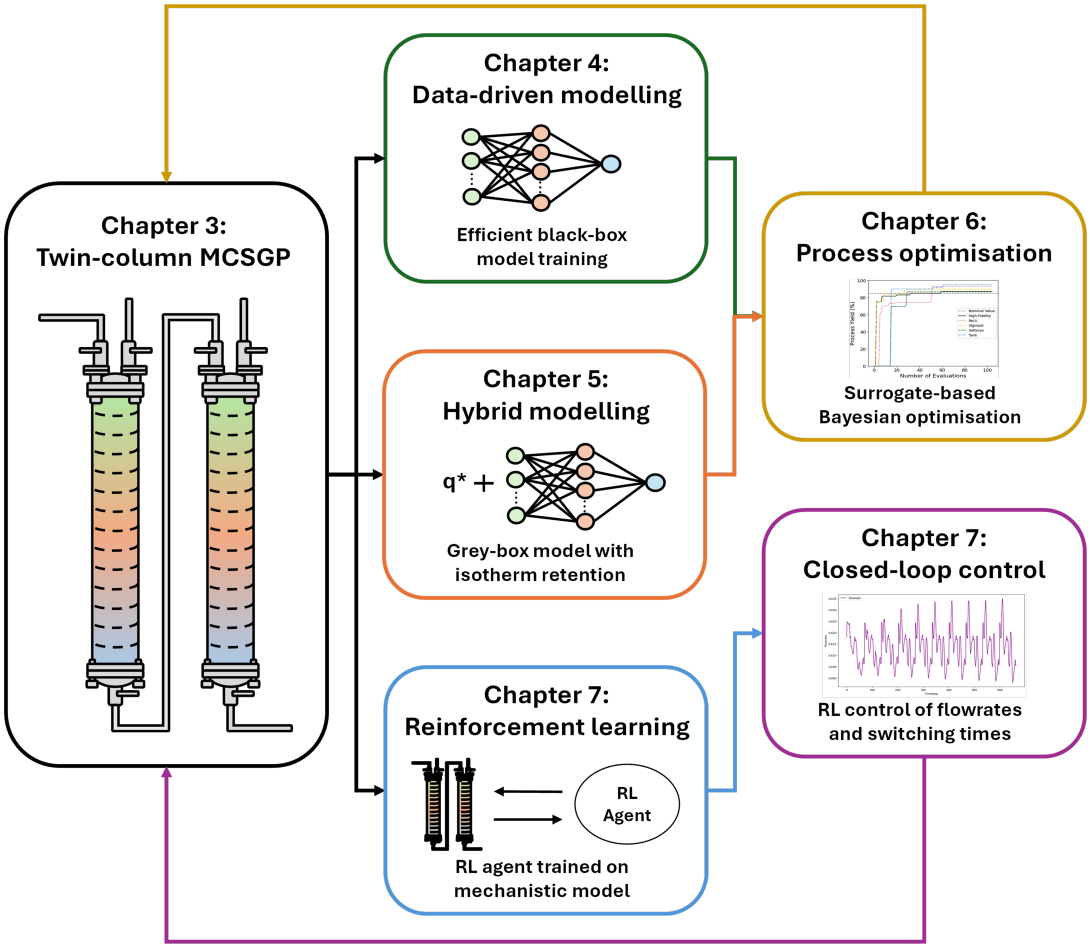


Figure 2.1: Thesis outline

Chapter 3

Case study

3.1 Introduction

The increasing demand for high-purity biologics such as monoclonal antibodies (mAbs), therapeutic peptides, and complex protein variants has highlighted the limitations of conventional batch chromatography, particularly in the presence of closely eluting impurities with a similar chromatographic profile to the target product (Luca et al., 2020). In such systems, the need to ensure product purity often requires the exclusion of overlapping regions in the chromatogram, inherently resulting in a yield–purity trade-off (Müller-Späth et al., 2013). This constraint has driven significant innovation in multicolumn chromatographic systems capable of continuous or semi-continuous operation with internal recycling of impure fractions. Among these, the multicolumn countercurrent solvent gradient purification (MCSGP) process has emerged as a prominent solution.

The MCSGP process was first introduced by Aumann and Morbidelli (2007) as a fully continuous countercurrent system capable of performing ternary separations under solvent gradient conditions. Unlike simulated moving bed (SMB) chromatography, which is well-suited for binary, isocratic separations, MCSGP was developed to address the challenges of centre-cut purifications involving weak impurities (W), the main product (P), and strong impurities (S), using linear solvent gradients (Krättli et al., 2013a).

The initial MCSGP configuration employed six chromatographic columns arranged to alternate between interconnected and batch phases (Aumann and Morbidelli, 2007). These early systems demonstrated significant gains in yield and productivity, particularly in the context of therapeutic peptide purification and mAb variant separation (Müller-Späth et al., 2008). However, the complexity and operational demands of such setups limited their broader industrial uptake. To address this, subsequent designs reduced the number of columns without sacrificing process performance. A three-column MCSGP system was introduced for semi-continuous operation, later extended to a four-column configuration to enable continuous feed loading and cleaning-in-place cycles (Krättli et al., 2011; Müller-Späth et al., 2013).

The most compact and industrially relevant format to date is the twin-column MCSGP, first described by Krättli et al. (2013b). This design utilises two identical columns that alternate between interconnected and batch modes, enabling internal recycling of overlapping W/P and P/S fractions within a single cy-

cle (Müller-Späth et al., 2013). The reduced hardware footprint and enhanced flexibility make this setup particularly well-suited for intensified biomanufacturing environments, where space, cost, and buffer usage are tightly constrained (Steinebach et al., 2017).

The effectiveness of the twin-column MCSGP has been demonstrated across a range of biologics. For instance, Müller-Späth et al. (2013) reported its application to the purification of bispecific antibodies, achieving high purity and yield while maintaining low solvent consumption. Similarly, Luca et al. (2020) applied the system to the purification of a therapeutic peptide, observing a fourfold increase in yield and tenfold improvement in productivity compared to an optimised batch process. These results underscore the process's ability to recover product from overlapping zones, thus reducing losses and minimising the need for multiple polishing steps.

In parallel to process design improvements, efforts have been made to enhance MCSGP process operation through model-based and data-driven control strategies. Early designs relied on empirical tuning based on preparative batch chromatograms, where gradient slopes, modifier concentrations, and switching times were selected to balance yield and purity (Grossmann et al., 2010). More recent studies implemented closed-loop control architectures, including model predictive control (MPC) and cycle-to-cycle repetitive control, which adjust operational parameters online based on chromatographic signals such as UV absorbance (Krättli et al., 2013b; Papathanasiou et al., 2016).

Given its cyclic operation, internal fraction recycling, and compatibility with gradient elution, the twin-column MCSGP process provides a well-defined and industrially relevant benchmark for the development and evaluation of advanced modelling and control strategies. In this Chapter, the system is introduced as the case study used to motivate the modelling, optimisation, and control methods developed in this work.

3.2 The twin-column MCSGP process

The analysis presented in this Chapter is focused on three main parts: (i) operating principles of the twin-column MCSGP process, (ii) mathematical formulation of the system, and (iii) key challenges associated with the system. Together, these provide the conceptual and methodological foundation for the remainder

of the thesis, as the assumptions, modelling choices, and optimisation and control strategies developed in subsequent chapters are directly informed by the characteristics and dynamics of the MCSGP process described here.

3.2.1 The process

The twin-column MCSGP process is a continuous ion-exchange chromatographic system in which two identical columns alternate between interconnected and batch modes (Krättli et al., 2013b). This configuration enables internal recycling of overlapping product and impurities fractions, allowing improved separation performance compared to conventional batch operation. A schematic representation of the cycle is shown in Figure 3.1.

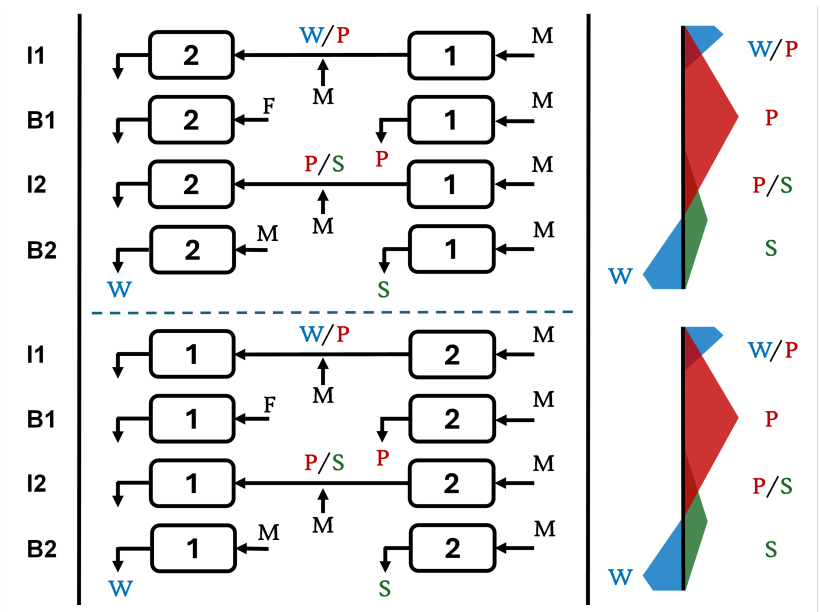


Figure 3.1: Schematic overview of a complete cycle of the twin-column multicolumn countercurrent solvent gradient purification process.

At process start-up, both columns are equilibrated with the initial buffer concentration and no feed has yet been introduced. The first cycle begins with phase I1, where a linear modifier (M) gradient is applied to the column at the right-hand side (column 1) and the flow is directed through this column to the one at the left-hand side (column 2). Since column 1 is empty at this point, no W/P fraction is available to recycle, and the phase primarily establishes the modifier gradient across the system. In phase B1, the columns are disconnected,

and fresh feed (F) is introduced for the first time into column 2, where separation begins under gradient conditions.

The cycle continues with phase I2, during which the columns operate again in an interconnected mode. As with phase I1, no P/S fraction is present, so only the modifier gradient continues through the column 1, while the separation of the mixture is executed in column 2. In phase B2, the weak impurities reach the outlet of column 2, and are therefore eluted and discarded, until the product begins to reach the outlet. At this point, the columns switch places and the full MCSGP cycle is initiated.

From the second half of the first process cycle onwards, the process follows a fixed, repeatable structure. During phase I1, the front of the elution profile from the right-hand column (column 2), containing a W/P mixture, is transferred to the left-hand column (column 1). Prior to entry, the recycled stream is diluted inline with a low-modifier buffer to promote re-adsorption of the product. Once the weakly adsorbing impurities stop being eluted from column 2 and a pure product stream is eluted, phase B1 commences and the interconnection between the two columns is interrupted. At this point, the pure eluted product is collected from the column 2 and fresh feed is introduced in column 1.

Once the strongly adsorbing impurities reach the outlet of column 2, and the product no longer meets the purity requirements for collection, the process switches to phase I2. In phase I2, the back end of the elution profile from column 2, containing a mixture of P/S, is internally recycled to column 1 after inline dilution. Finally, once the product stops being eluted, the process transitions to phase B2, where the weak impurities are discarded from column 1 and strong impurities from column 2, signifying the end of one process cycle. The columns then switch positions, and the sequence repeats. Through successive cycles, the system approaches cyclic steady state (CSS), where concentration profiles stabilise and are identical from one cycle to the next.

3.2.2 Mathematical modelling

The mechanistic model used for the simulation of the process has been developed and experimentally validated by Müller-Späth et al. (2008). The model parameters, including axial dispersion, mass transfer, and adsorption isotherm coefficients, were obtained from single-column gradient elution experiments and

applied directly to the simulation of twin-column system without further adjustment. These parameters are not described here due to confidentiality reasons. The model accurately reproduced experimental elution profiles and key performance indicators (KPIs), supporting its use for the twin-column MCSGP system. Each column is simulated independently, with the inlet and outlet conditions updated dynamically according to the current process phase. During interconnected phases, the outlet of one column becomes the inlet of the other, whereas in batch phases the columns operate in isolation with boundary conditions defined by the feed or eluent streams. Switching events determine when and how the columns are interconnected, whereas dilution of the recycled fractions is modelled by adjusting the modifier concentration at the column inlet during the interconnected phases. The main equations describing the system are presented below.

Liquid phase concentration of components:

$$\frac{\partial c(z,t)_{i,h}}{\partial t} = D_{ax}\frac{\partial^2 c(z,t)_{i,h}}{\partial z^2} - \frac{Q_h}{A_{col}\varepsilon_i}\frac{\partial c(z,t)_{i,h}}{\partial z} - \frac{(1-\varepsilon_i)}{\varepsilon_i}\frac{\partial q(z,t)_{i,h}}{\partial t} \tag{3.1}$$

where, t is the time, z the column length, i the components in order of elution (1. modifier, 2. weak impurities, 3. product, and 4. strong impurities), h the column index, $c(z,t)_{i,h}$ the liquid phase concentration of component i in column h, $q(z,t)_{i,h}$ the solid phase concentration of component i in column h, A the column cross-section, Q_h the flow rate of column h, ε_i the column porosity for component i, and D_{ax} the axial dispersion coefficient. The concentration of the liquid phase for the components is dependent on the dispersion ($D_{ax}\frac{\partial^2 c(z,t)_{i,h}}{\partial z^2}$), the convective transport ($\frac{Q_h}{A_{col}\varepsilon_i}\frac{\partial c(z,t)_{i,h}}{\partial z}$), and the adsorptive interaction ($\frac{(1-\varepsilon_i)}{\varepsilon_i}\frac{\partial q(z,t)_{i,h}}{\partial t}$) of each component i within the column.

Solid phase concentration of components:

$$\frac{\partial q(z,t)_{i,h}}{\partial t} = k_i\left(q^*\left(c(z,t)_{i,h}\right) - q(z,t)_{i,h}\right) \tag{3.2}$$

where, $q^*(c(z,t)_{i,h})$ is the equilibrium solid phase concentration of component i in column h, and k_i the lumped mass transfer coefficient of component i. This first-order formulation combines film and pore diffusion into a single rate parameter. In the twin-column MCSGP context, mass transfer kinetics strongly influence the efficiency of re-adsorption when impure product fractions are recycled in phases I1 and I2.

Initial conditions:

$$c(z, t = 0)_{i,h} = c^0_{i,h} \tag{3.3}$$

where, $c^0_{i,h}$ is the initial liquid phase concentration of component i in column h at the start of the process operation.

Boundary conditions:

$$c(z = 0, t)_{i,h} = c^{in}_{i,h} \tag{3.4}$$

$$\frac{\partial c\left(z = L_{col}, t\right)_{i,h}}{\partial z} = 0 \tag{3.5}$$

where, $c^{in}_{i,h}$ is the inlet liquid phase concentration of component i in column h. The former boundary condition (Equation 3.4) indicates that at any time point during the process at the beginning of each column the concentrations are identical to the inlet concentrations for each component (either coming from the other column or the feed itself depending on the mode of operation). The latter (Equation 3.5) indicates that at the end of each column the liquid phase concentrations of the species depend solely on the time point of the process and do not change throughout space. This means that the liquid phase concentrations at the end of the column are the concentrations reported during the collection points of the separation species, as well as the concentrations used as inputs when the columns are interconnected.

Competitive Bi-Langmuir isotherm – Equilibrium solid phase concentration of separation components:

$$q^*\left(c(z,t)_{i,h}\right) = \frac{c_{i,h} \cdot H^I_{i,h}}{1 + \sum_{i=2}^{n_{\text{comp}}} \frac{c_{i,h} \cdot H^I_{i,h}}{q^I_{i,h}}} + \frac{c_{i,h} \cdot H^{II}_{i,h}}{1 + \sum_{i=2}^{n_{\text{comp}}} \frac{c_{i,h} \cdot H^{II}_{i,h}}{q^{II}_{i,h}}} \tag{3.6}$$

where, $q^I_{i,h}$ and $q^{II}_{i,h}$ are the saturation capacities for the adsorption sites 1 and 2 of component i in column h, and $H^I_{i,h}$ and $H^{II}_{i,h}$ the Henry constants for the adsorption sites 1 and 2 of component i in column h, respectively. The competitive formulation of the separation isotherm captures the competitive adsorption behaviour between the separation species, namely the weak impurities, the product, and the strong impurities.

Equilibrium solid phase concentration of modifier:

$$q^* \left(c(z,t)_{1,h}\right) = H_{mod} \cdot c(z,t)_{1,h} \tag{3.7}$$

where, H_{mod} is the Henry constant of the modifier. Contrary to the separation species, the equilibrium solid phase concentration of the modifier is only a linear function of its liquid phase concentration.

Henry constants and saturation capacities for product and impurities:

$$H_{i,h} = a_{i,1} \cdot \left(c(z,t)_{1,h}\right)^{a_{i,2}} \tag{3.8}$$

$$H_{i,h}^{II} = a_{i,3} \cdot \left(H_{i,h}\right)^{a_{i,4}} \tag{3.9}$$

$$H_{i,h}^{I} = H_{i,h} - H_{i,h}^{II} \tag{3.10}$$

$$q_{i,h}^{I} = a_{i,5} \cdot \left(H_{i,h}^{I}\right)^{a_{i,6}} \tag{3.11}$$

$$q_{i,h}^{II} = a_{i,7} \cdot \left(H_{i,h}^{II}\right)^{a_{i,8}} \tag{3.12}$$

where, $H_{i,h}$ is the overall Henry constant of component i in column h, and $a_{i,1}$ to $a_{i,8}$ are experimentally defined species-specific coefficients.

Cyclic steady state convergence criterion:

$$|c(z = L_{col}, t)_{i,h}^{(j)} - c(z = L_{col}, t)_{i,h}^{(j-1)}| < \varepsilon_{\text{tol}} \tag{3.13}$$

where $c(z = L_{col}, t)_{i,h}^{(j)}$ is the outlet concentration of species i eluted from column h in cycle j at cycle timepoint t, and ε_{tol} the chosen tolerance.

Key performance indicators - product purity and process yield:

$$Pur_j = \frac{c_{\text{av}P,j}}{c_{\text{av}W,j} + c_{\text{av}P,j} + c_{\text{av}S,j}} \tag{3.14}$$

$$Y_j = \frac{c_{\text{av}P,j} \cdot Q_{right}}{c_{\text{feed}P} \cdot Q_{left}} \tag{3.15}$$

where, Pur_j and Y_j correspond to the average purity and recovery yield over a process cycle j, $c_{\text{av}i,j}$ is the average concentration of component i collected during cycle j, $c_{\text{feed}P}$ is the feed concentration of the product, and Q_{right} and Q_{left} are the flowrates applied to the right- and left-hand column respectively during product collection. Purity is calculated as the fraction of the product concentration in the collected stream, and yield as the ratio of recovered product mass from the right-hand column to that introduced in the feed in the left-hand column. Both KPIs are calculated over the product collection windows of the cycle, which occur during the B1 phases on the right-hand column.

3.2.3 System analysis

The twin-column MCSGP process is characterised by the direct coupling of the two chromatographic columns through material recycling, non-linear component interactions, and operating conditions that change abruptly at switching points. Unlike conventional batch chromatography, where disturbances remain confined to a single cycle, in the MCSGP configuration any fluctuation is immediately propagated through internal recycling and may continue to affect subsequent cycles until CSS is reached. This interdependence means that separation performance cannot be understood or optimised on the basis of individual column behaviour alone. Instead, the dynamic environment is shaped by the cumulative interaction of gradient-driven adsorption, fraction recycling, and abrupt phase changes, creating a system that is both operationally efficient and highly sensitive to parameter variation. Understanding these dynamics is essential to be able to accurately model the system and for the development of optimisation and control strategies that can operate reliably in real-time.

During the interconnected phases (I1 and I2), the outlet of one column serves directly as the inlet of the other, so the quality of the recycled fraction determines not only the recovery of the receiving column but also the stability of the overall process. Any deviation in modifier concentration or component ratio is therefore propagated from one column to the other, altering the progression of the separation fronts and shifting the balance between the separation species. Switching events further complicate the dynamics, as each switch resets inlet conditions, changing both the composition and modifier concentration of the feed to the receiving column. These transitions generate transient effects that propagate along the column and can alter the separation boundaries between weak impurities, product, and strong impurities.

The precise timing of switching events is critical as well, as early or delayed switching can either truncate the recovery of pure product or allow impurities to enter the product fraction, directly affecting yield and purity and therefore failing to meet regulatory specifications. Optimal switching points depend on the progression of component fronts within each cycle and can shift under variations in feed composition, flowrates, or gradient conditions. Since the process is inherently cyclic, maintaining optimal switching is essential for reaching and sustaining CSS, where concentration profiles at a given phase and position in the

column repeat identically from one cycle to the next.

When switching times are predetermined and fixed during operation, the performance of the twin-column MCSGP system is sensitive to multiple operating conditions (Figure 3.2). Changes in feed concentration (Figure 3.2b–c) primarily affect the magnitude and shape of the elution peaks. At reduced feed concentration, product and impurity peaks are significantly diminished, and their profiles broaden slightly, lowering overall recovery within the fixed collection window. At increased feed concentration, the concentration of the eluted product increases as well, but the higher component load also enhances competitive interactions and column saturation, leading to slight peak distortions. Variations in feed composition, whether in product concentration or impurity content, can shift the elution windows and move W/P or P/S boundaries into or out of the product collection window. Even small deviations can lead to off-spec product or reduced yield, making feed variability a key disturbance to manage.

Flowrate variations (Figure 3.2d–e) have a more pronounced effect on the separation and peak shape of the separation species. At reduced flowrates, the increased residence time leads to delayed elution and compression of the species' elution into a narrow, high-intensity peak. This can shift the product, along with the impurities, outside the predefined collection window, creating a risk of product loss under fixed cycle timing. In contrast, at elevated flowrate values, mass-transfer limitations dominate, as the interaction of the species with the column is reduced. As such, the product elution profile broadens substantially, overlaps with weak and strong impurities, and develops a distorted profile. This broadening reduces the effective selectivity of the process and increases the chance of impurities being co-eluted and therefore collected with the product.

Finally, similarly to the effect of the flowrate, the initial modifier concentration exerts a pronounced effect on the relative positioning and overlap of species (Figure 3.2f–g). A lower modifier concentration delays the onset of elution, creating sharp, delayed peaks for the product and the impurities elution profiles, whereas a higher modifier concentration accelerates elution, shifting the profiles earlier, while still increasing their degree of overlap. Both scenarios reduce the robustness of the process to fixed switching decisions, as the separation boundaries no longer align with the pre-determined cycle structure.

From a modelling perspective, the twin-column MCSGP process presents several distinct challenges. The first arises from the non-linear nature of competitive

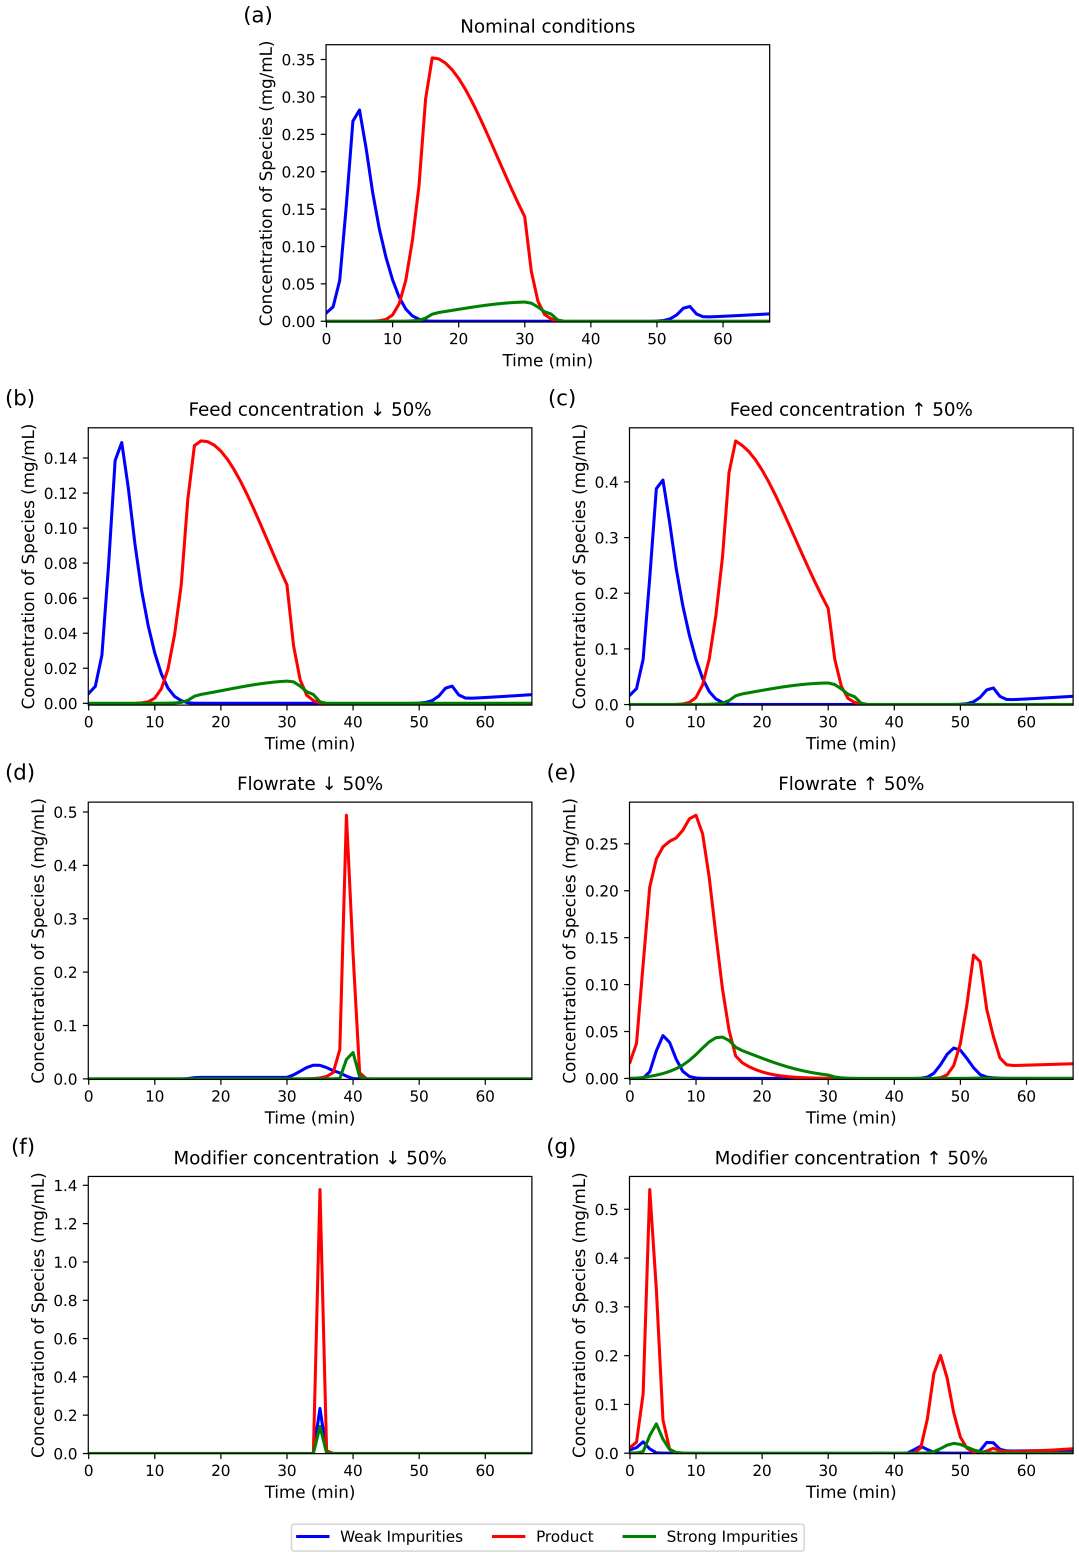


Figure 3.2: Effect of operating condition variations on elution profiles at CSS. Each case shows one cycle with a ±50% change from nominal operating conditions in feed concentration, flowrate, or initial modifier concentration.

adsorption under dynamically changing modifier concentrations, which governs the progression and overlap of all separation fronts. Capturing this behaviour requires models that can account for strong interactions between species and rapidly shifting adsorption equilibria.

A second challenge comes from the presence of extended zero-concentration regions for many states during parts of the cycle (Figure 3.3). In mechanistic simulations, these regions introduce stiffness into the numerical solution, slowing convergence. For data-driven representations, they lead to sparsity in the training distribution, since most datapoints correspond to trivial zero outputs while only a small subset captures informative peak dynamics. This imbalance reduces the effective information content of the dataset, increases the risk of overfitting to high-concentration regions or learning degenerate solutions that minimise error by predicting zeros, and complicates the ability of models to generalise across phases.

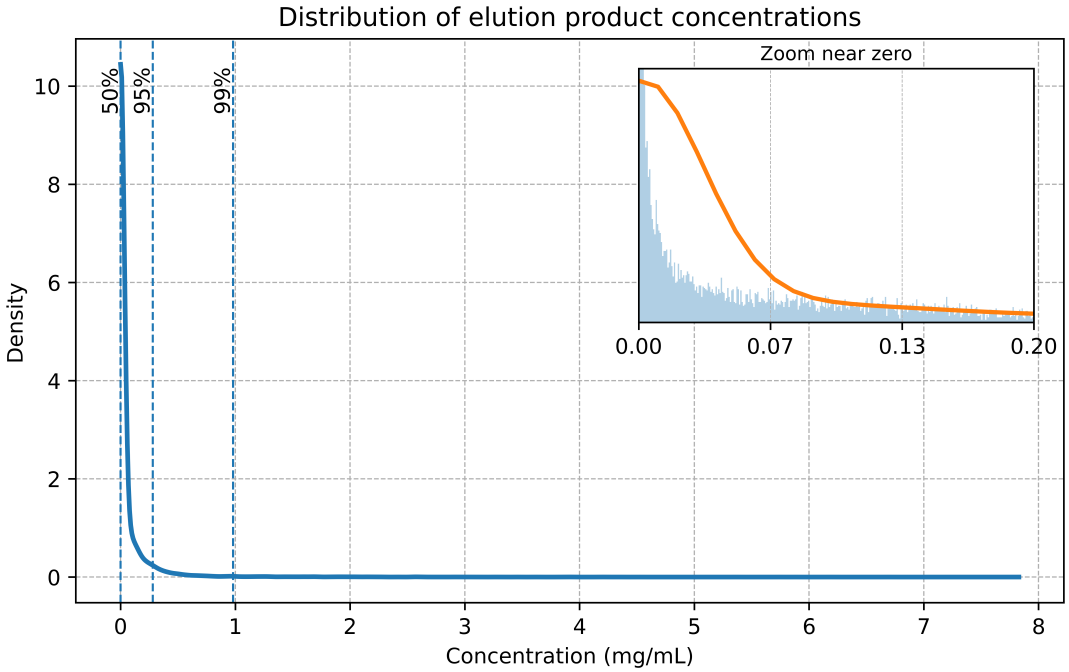


Figure 3.3: Distribution of elution product concentration values at CSS across 400 process cycles.

A further complication is introduced by the phase-dependent connectivity of the columns, which changes inlet conditions abruptly at switching points. Each switch resets both composition and modifier concentration, creating sharp transients that propagate along the column and must be handled with high

numerical accuracy to preserve mass balance. Since these transients repeat with every cycle, their effect accumulates until CSS is reached.

Finally, the convergence to CSS itself can be slow, requiring more than ten full process cycles (Papathanasiou et al., 2016). This significantly increases the computational burden of mechanistic simulation, particularly in cases where many operating points or disturbance scenarios must be explored. When combined, these features of non-linear dynamics, sparse and imbalanced data distributions, abrupt switching events, and long convergence horizons make direct mechanistic optimisation resource-intensive and highlight the need for reduced-complexity representations that can be used in online settings.

3.3 Remarks

The analysis presented in this Chapter has outlined the operating principles of the twin-column MCSGP configuration, provided the mechanistic framework used for its simulation, and identified the key sensitivities and non-linearities that govern its performance. The system's dynamic complexity, strong column coupling, and strict purity requirements create a need for advanced process control and optimisation methods. Robust operation demands strategies capable of handling variable and uncertain feed composition, the non-linear effects of gradient elution, and the transient behaviour introduced by switching. These characteristics make the process an industrially relevant and technically demanding benchmark for the development of advanced modelling, optimisation, and control strategies.

The methodologies presented in the subsequent Chapters build directly on this foundation. Chapters 4 and 5 develop and evaluate data-driven and hybrid modelling approaches for representing the twin-column MCSGP process with reduced computational complexity. Chapter 6 applies these models within a surrogate-based optimisation framework to explore operating parameter selection under realistic constraints. Finally, Chapter 7 introduces and assesses reinforcement learning-based control strategies capable of managing the process in real time under disturbance conditions. Together, these methods aim to address the challenges identified here and to provide a general framework for advanced decision-making in complex chromatographic systems.

Chapter 4

Data-driven modelling

Outputs of this chapter have been disseminated in:

- **F. Michalopoulou** and M. M. Papathanasiou. Assessment of data-driven modeling approaches for chromatographic separation processes. *AIChE Journal*, e18600, 2024. (Invited article in the *2024 Futures Issue*)

4.1 Introduction

Chromatographic separation processes are widely used in the purification of complex biological products, particularly in biopharmaceutical manufacturing (Strube et al., 2012). As purity specifications tighten, productivity targets rise, and development timelines shorten, the demand for real-time optimisation and control increases. Meeting these requirements calls for modelling approaches that balance predictive accuracy with computational efficiency (Rathore et al., 2023; Yang et al., 2023). First-principles models have long served as the foundation for the simulation of chromatographic separation systems, providing a physically meaningful description of convection, axial dispersion, and adsorption phenomena (Chen et al., 2020). Such models are often computationally intensive and rely on spatial and temporal discretisation schemes that can result in thousands of equations to be solved per model iteration, especially in continuous multicolumn setups, that operate under cyclic steady state (CSS).

The complexity of mechanistic models can pose a significant challenge when simulations must complete on decision-making timescales, often within seconds to minutes, as required for digital twins, online optimisation, and real-time control (Kawajiri, 2021). In such contexts, even a single model evaluation can take several minutes, and the time required to perform tasks online becomes prohibitive. Moreover, the development of these models is highly dependent on expert knowledge and extensive experimental data for parameter calibration, often making their construction costly and time-consuming (Qu et al., 2024; Shekhawat and Rathore, 2019).

To address these challenges, surrogate modelling approaches have been proposed as a way to approximate system behaviour while reducing the computational burden (Osberghaus et al., 2012; Solle et al., 2017). These models are typically trained on simulation and/or experimental data and aim to replicate the outputs of high-fidelity models without requiring a full mechanistic formulation. In the field of chromatography, several types of surrogates have been explored, ranging from polynomial regression and response surface models to Gaussian processes and machine learning techniques (Ding and Ierapetritou, 2024; Kalil et al., 2000; Tang et al., 2023; Subraveti et al., 2022b). Among these, artificial neural networks (ANNs) have attracted particular attention owing to their flexibility in representing non-linear relationships and their capacity to scale to

high-dimensional problems.

Although neural network surrogates have been applied to chromatographic problems before, their adoption has often been limited to specific components of the simulation pipeline, for instance, the prediction of single elution peaks, local isotherms, or process key performance indicators (KPIs), such as product purity and process yield (Mouellef et al., 2023; Oh et al., 2007; Wang et al., 2017a). Few studies have evaluated their performance in capturing full-cycle outlet concentration profiles in cyclic systems, particularly when the input space includes a broad range of operating conditions (Li and Wang, 2024, 2025). Moreover, when ANNs are used in process applications, there is often limited discussion of model structure, input feature selection, or training data requirements, despite these factors having a significant impact on the reliability and generalisability of the models.

To address these gaps, this Chapter investigates the development of pseudo-dynamic ANN-based data-driven models trained on in silico generated mechanistic data for the twin-column multicolumn countercurrent solvent gradient purification (MCSGP) process, operating under CSS conditions. The objective is to assess to what extent such surrogates can predict outlet concentration profiles across a cycle without the need for spatial discretisation or direct time inputs, so as to enable rapid model simulation. Two model configurations are considered: one where each separation species' outlet concentration is predicted separately using dedicated ANNs, and one where all concentrations are predicted jointly by a single ANN. Both rely on a shared input feature structure that includes process flowrates, inlet concentrations, and modifier gradients.

The focus is on evaluating the models' ability to predict the concentration elution profiles of the separation species within the models' training bounds and with reduced time requirements. In contrast to approaches that target system KPIs, this approach allows for the prediction of the elution concentrations of the separation species at each time point throughout the process, thus facilitating the models' integration in process optimisation frameworks, where decisions need to be made at short time intervals. The results are presented with an emphasis on the models' predictive accuracy, structure, and computational performance, laying the groundwork for later chapters that build on these surrogates for hybrid modelling and optimisation.

4.2 Methodology

The data-driven models presented in this Chapter are developed based on four main steps: (i) analysis of input–output structures, (ii) synthetic data generation, (iii) identification of candidate model architectures, and (iv) model training and hyperparameter tuning. The methodologies underlying each of these steps are discussed in the following sections.

4.2.1 Input-output structure

The data-driven models developed in this Chapter aim to facilitate rapid online applications, such as process optimisation and control of the twin-column MCSGP process. Compared to the mechanistic model of the system, these models aim to eliminate the main computational bottlenecks, so as to reduce computational cost and simulation time. Based on the system analysis presented in Chapter 3, the main sources of computational complexity presented in the model are centred around the partial differential and algebraic equations (PDAEs) that describe it, which demand spatial discretisation of the model and therefore greatly increase the number of variable and equation sets that need to be solved per model iteration. Given the cyclic operation of the process, this computational overhead is further amplified by the multiple cycle simulations that are required in order to reach CSS.

To address these bottlenecks, a data-driven modelling approach is implemented, in order to predict directly the liquid phase concentrations of the separation species being eluted at the column outlets at CSS. The first step in developing such a model is defining an input–output structure that reflects the system's underlying dynamics while remaining computationally tractable. For this purpose, an assessment of how each operating condition influences the concentration profiles of the mixture components over a cycle is performed. In the twin-column MCSGP system, component retention and elution are governed by a combination of flowrates, modifier concentration gradients, and feed composition. These variables not only affect the system independently but also interact across phases through the interconnected operation of the two columns.

As demonstrated by Papathanasiou et al. (2016), the modifier concentration at the start of each phase plays a central role in controlling the interactions

both between the separation species themselves, and between the species and the solid phase of the chromatography columns. As such, even small variations in the modifier concentration can shift the relative retention times and elution profiles of the closely eluting species. Since the modifier concentration in the MCSGP system is typically adjusted at the beginning of each phase and then a steady gradient is employed, the initial values for both columns are included as inputs to capture these transitions.

Likewise, the flowrates applied during each phase affect the residence time of the species and determine how much time they have to interact with the solid phase. Higher flowrates shorten contact with the stationary phase and reduce retention, leading to earlier elution and sharper fronts, whereas lower flowrates prolong contact, increase retention, and can delay and broaden the elution peaks (Krättli et al., 2013a). Including the flowrates of each phase in the model input allows the surrogate to account for these effects and adapt the predicted outlet trajectories accordingly.

In addition to the modifier concentration and the flowrates, the composition of the feed stream also affects the shape and positioning of elution peaks. Higher product concentration can result in earlier breakthrough and tailing, while changes in impurity levels can alter the overlap between target and non-target components (Fioretti et al., 2024). In the case at hand, the twin-column MCSGP system is used as the final purification step for the downstream processing of a monoclonal antibody (mAb). In that context, the feed composition is not fixed but reflects the output of the previous downstream separation step. For this reason, it is treated as a measured disturbance, with an expected range of variation defined by upstream processing. It is therefore included as an input to ensure that the data-driven models are able to adapt to changes in the feed composition.

Once the input variables are selected, the process outputs of interest are defined. In this work, the models are trained to predict the liquid phase concentrations of the weak impurities (W), product (P), and strong impurities (S) at the column outlet. These are the same concentrations used in the mechanistic model to compute product purity (Equation 3.14) and process yield (Equation 3.15), the two process KPIs of interest. Both KPIs are calculated over the product collection windows of the cycle, which occur during the B1 phases on the right-hand column. The advantage of predicting the concentrations of the species at

each time step of the process, rather than directly predicting the KPIs, lies in the ability to simulate the process dynamically. This is particularly important in twin-column MCSGP, where the two columns operate in an interconnected manner. As such, the elution profiles of the species can be reconstructed, providing further insight into the performance of the process.

4.2.2 Synthetic data generation

The experimentally validated mechanistic model of the twin-column MCSGP process(Chapter 3) is used as a virtual representation of the system to generate synthetic data. Based on the previously identified input-output structure, the operating conditions presented in Table 4.1, are varied within the validation range of the mechanistic model, using a quasi-random Sobol sequence.

Table 4.1: Data generation inputs.

Input	**Lower bound**	**Upper bound**
Feed concentrations		
Feed concentration of weak impurities (mg/mL), ($C_{\text{feed}W}$)	0.05	0.08
Feed concentration of product (mg/mL), ($C_{\text{feed}P}$)	0.3	0.5
Feed concentration of strong impurities (mg/mL), ($C_{\text{feed}S}$)	0.03	0.05
Flowrates		
Flowrate during phases B1 and B2 (mL/min), (Q_B)	0.1	1
Flowrate during phases I1 (mL/min), (Q_{I1})	0.1	1
Flowrate during phases I2 (mL/min), (Q_{I2})	0.1	1
Modifier concentrations		
Feed concentration of modifier (mg/mL), ($C_{\text{feed}M}$)	1	3
Initial modifier concentration of Phases I1 for column 1 (mg/mL), (C_{M1})	2	3
Initial modifier concentration of Phases I1 for column 2 (mg/mL), (C_{M2})	0.8	1.2

Chromatographic separation is highly influenced by the timing and gradient of the product elution profile. Consequently, it is essential that the data-driven models incorporate temporal information about the concentrations of mixture components to effectively capture the process dynamics. For this reason, the input set considered for model training is expanded to include the inlet concentrations of all three mixture components at each time step, along with the corresponding modifier concentration. This addition aids in capturing the dynamic nature of the system, by capturing the effect of changes in the inlet concentrations on the tracked outputs.

This pseudo-dynamic input strategy is implemented by sampling the inlet

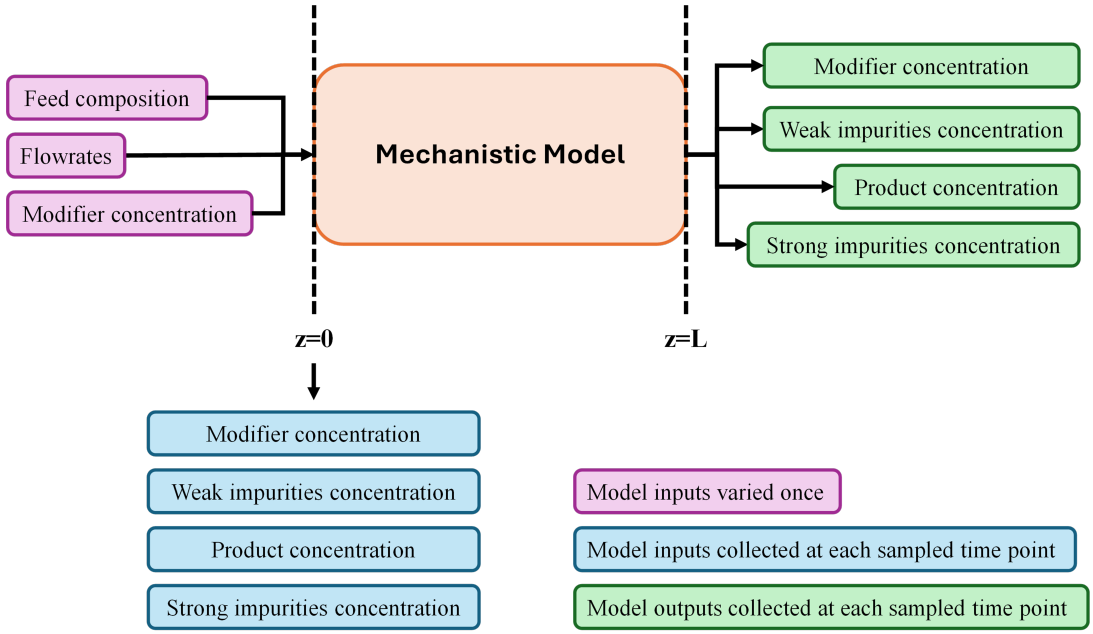


Figure 4.1: Schematic representation of the input–output structure for the synthetic data generation.

and outlet concentrations at regular time intervals ($dt = 1\ min$), as outlined in Figure 4.1. Simultaneously, the data-driven models receive the time-varying modifier concentration (i.e. the gradient profile) as an input, so that changes in the gradient are explicitly reflected in the predicted elution behaviour. This achieves a representation of temporal evolution through measured trajectories rather than an explicit time variable, so that predicted elution profiles are linked to the inlet concentrations and operating conditions rather than tied to a fixed timeline. This offers three key advantages: relying solely on measurable variables, reducing computational overhead, and enabling the model to accommodate different switching times. The latter is particularly beneficial during process development stages, where switching times are yet to be optimised and fixed.

Importantly, the input–output structure presented here is defined entirely within the mechanistic modelling framework and does not rely on process measurements. All inputs are known *a priori* operating variables, including the feed composition, mixture component inlet concentrations, flowrates, and modifier concentrations, allowing the data-driven models to be evaluated independently of data acquisition and in parallel with process operation. The model is trained to predict the outlet concentrations of the three components, along with the modifier, which is included because it is recycled between the two chromatographic

columns and therefore affects the separation being executed at each time on the left-hand column of the system. This departs from existing approaches that rely on the prediction of KPIs, which are typically determined offline in a separate analytical step, limiting their ability to retain detailed temporal information about the process. By predicting complete outlet trajectories directly, the proposed structure preserves the dynamic information of the separation process and avoids reliance on reduced representations that obscure temporal behaviour.

For the purposes of this work, both the modifier concentration gradient and the column length are fixed at their nominal values, consistent with the validation range of the mechanistic model, while the duration of each step is held constant, and no start-up or shutdown scenarios are considered. For each input configuration, the mechanistic model is simulated until CSS is achieved, and the corresponding inputs and outputs are recorded for one process cycle. Based on the three identified groups of manipulated variables, namely the feed concentrations, flowrates, and modifier concentrations (Table 4.1), input combinations are designed to form four distinct subsets. These include three subsets generated by independently varying each group of inputs, and one subset where all inputs are varied simultaneously. This sampling strategy allows the resulting dataset to capture both individual and combined effects of the input variables on the process outputs.

Further to that and given the limited availability of experimental data in biomanufacturing, it is important to evaluate the feasibility of developing accurate data-driven models using relatively small datasets. Four datasets comprising 50, 100, 200, and 400 process cycles are therefore generated to assess the impact of dataset size on model performance. This provides insight into the quantity of data required for the development of reliable models in data-constrained settings, particularly where mechanistic models may not be accessible.

4.2.3 Model architecture

The development of data-driven models for continuous chromatographic separation processes is challenged by the intrinsic non-linearity, periodicity, and dynamic behaviour of the system. These characteristics, coupled with strong interdependencies among input and output variables, necessitate the careful selection of model architecture and internal components such as activation functions.

In ANNs, the activation function governs the transformation applied to the weighted input of each neuron and is central in determining the model's capacity to approximate complex relationships. To ensure a comprehensive evaluation of model performance, four activation functions are considered, (i) ReLU, (ii) Sigmoid, (iii) Softmax, and (iv) Tanh, each representing distinct mathematical characteristics in terms of non-linearity, smoothness, and saturation behaviour (Dubey et al., 2022). This selection of activation functions aims to capture a broad spectrum of functional responses that may influence the ANN's ability to approximate the dynamic and periodic behaviour of chromatographic systems. Given the absence of established *good-practices* for data-driven modelling in this domain, a systematic comparison is required to identify which functions are best suited to represent the underlying process dynamics and variable interactions.

To this end, two distinct ANN-based data-driven modelling strategies are considered, based on the input structure identified, consisting of 9 operating conditions and 4 inlet concentrations. The first is a multiple-input/multiple-output (MIMO) approach, wherein a single multi-output network with a 13×4 structure is developed, designed to simultaneously predict all four process outputs from the same input vector (Figure 4.2a). The second, a multiple-input/single-output (MISO) approach, consists of four independent 13×1 networks, each trained to model one output concentration in isolation (Figure 4.2b). This architectural distinction allows for a comparative evaluation of shared versus specialised representation learning in the context of strongly coupled output behaviours. For both approaches, the four activation functions are systematically tested to assess their ability to support convergence, mitigate vanishing gradient issues, and ultimately enable accurate generalisation across the input domain.

4.2.4 Hyperparameter tuning

Bayesian optimisation (BO) is employed to tune the hyperparameters of the ANNs, namely the number of hidden layers, the number of neurons per layer, and the learning rate (Table 4.2). This method is selected over conventional grid search due to the size and complexity of the training dataset, which renders exhaustive search approaches inefficient and computationally prohibitive. The optimisation is guided by the mean squared error (MSE) calculated on the testing subset, as it quantifies the deviation between predicted and actual concentrations.

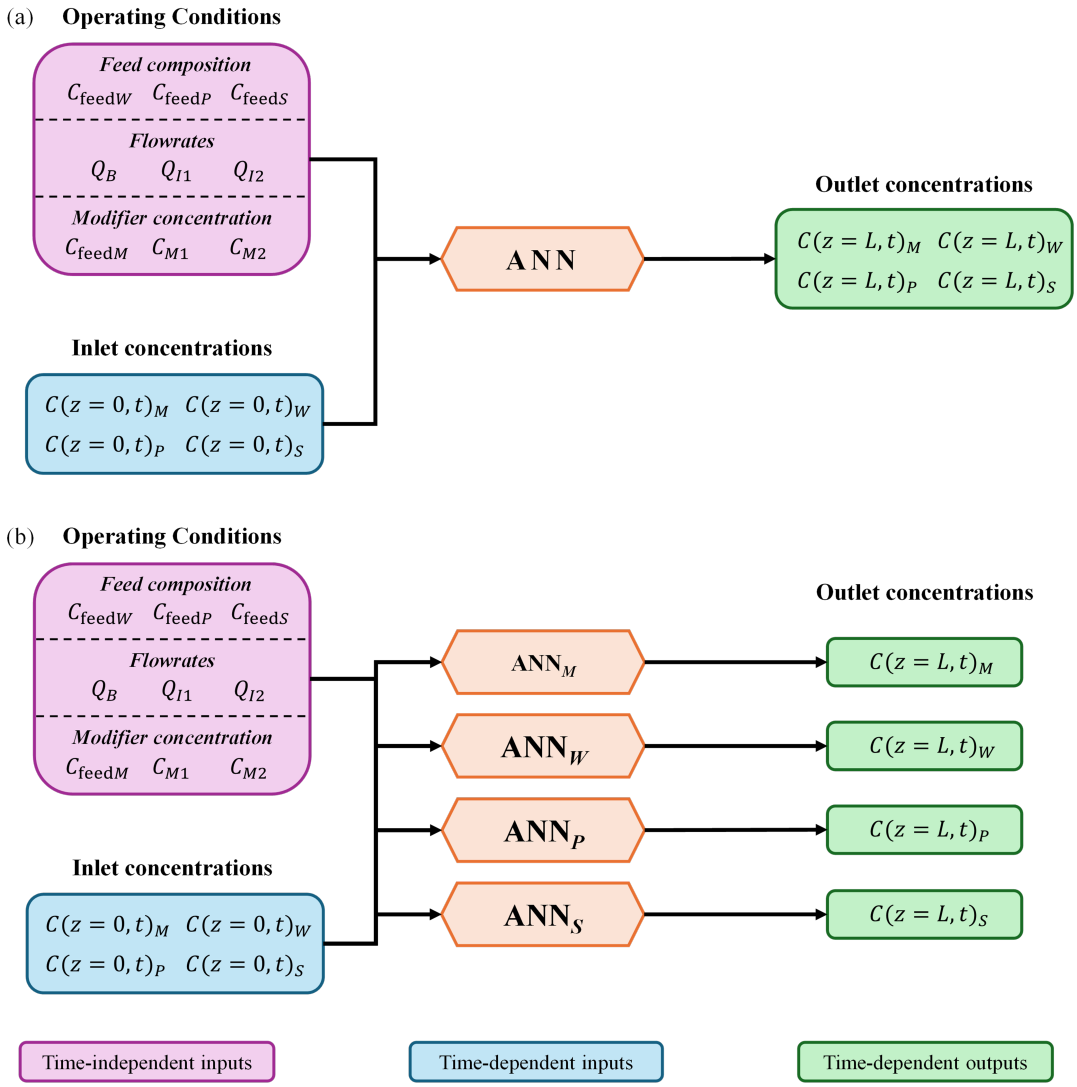


Figure 4.2: (a) Schematic representation of the structure of the data-driven models using MIMO ANNs, (b) Schematic representation of the structure of the data-driven models using MISO ANNs.

Table 4.2: Hyperparameter bounds for the Bayesian optimisation.

Hyperparameters	**Lower bound**	**Upper bound**
Hidden layers	1	5
Neurons per hidden layer	16	128
Learning rate	10^{-5}	0.1

All architectures are trained using the Adam optimisation algorithm, while the training is performed using a fixed split of 80% training and 20% testing data, with a maximum of 40,000 epochs. To avoid overfitting, early stopping is implemented if the model's performance on the testing dataset fails to improve over 500 consecutive epochs. Regardless of the activation function used in the hidden layers, the output layer employs a linear activation function to preserve continuity and avoid saturation effects in the final regression outputs. The number of neurons is kept constant across all hidden layers to reduce the dimensionality of the hyperparameter space and simplify the optimisation task.

Following the tuning of the hyperparameters, the best-performing ANN is selected for each combination of dataset size, activation function, and input–output configuration. These models are then integrated into the process schedule of the twin-column MCSGP chromatography system and evaluated based on their predictive accuracy and computational cost. When integrated into the process schedule, the ANNs operate as unconstrained surrogates of the mechanistic system. As a result, mass conservation is not explicitly enforced, and exact mass balance closure is not guaranteed for the predicted concentration trajectories. As both columns in the system are identical in terms of design and physicochemical characteristics (Figure 3.1), the same ANN architecture is applied to both.

4.3 Results & discussion

4.3.1 The multiple-input/multiple-output approach

Achieving high product purity is the principal objective in bioprocess purification, with yield serving as a secondary metric reflecting overall process efficiency. As previously described, KPIs are calculated at the end of each process cycle, based on the average outlet concentrations during the product collection window

(Equations 3.14 and 3.15). It is therefore essential that any data-driven model is capable of accurately capturing the system dynamics within this specific time frame, in addition to the prediction of the entire elution profile.

In the ternary mixture considered here, the product constitutes the critical component to be recovered, whereas the impurities are to be removed. For this reason, the results presented below focus initially on the model's predictive performance for the product concentration during the elution window. Predictions for the weak and strong impurities are also presented to provide a complete assessment of model performance, as both must be tightly controlled during operation and maintained below specified thresholds to satisfy purity constraints. Information regarding the optimal structure of the ANNs integrated in the models, as identified through hyperparameter tuning, can be found in Appendix A.

4.3.1.1 Product

First, the performance of the data-driven models developed using the MIMO approach is examined. In this approach, a single network is trained to predict the elution profiles of the three mixture components and the modifier simultaneously. The influence of four different activation functions and four dataset sizes is assessed. Across all cases, the data-driven models exhibit some degree of deviation from the reference profiles, typically under- or overestimating concentrations at different points within the elution window.

These discrepancies can be attributed to the network structure itself, which is designed to learn shared representations in the common hidden layers used for the prediction of all outputs. Because the outputs are predicted concurrently using the same set of internal layers and parameters, the ANNs inherently introduce correlations among the elution profiles. While this structure allows the models to reflect the coupled behaviour observed in chromatographic separations, it also implies that errors in the prediction of one component may propagate to the others. This interdependency can limit the models' ability to independently predict each output, particularly in regions of the profile where component behaviour diverges.

As illustrated in Figure 4.3, there exists a clear trend between the number of training samples and the performance of the data-driven models, as model performance improves as the size of the training dataset increases. The ReLU-

activated models exhibit the most stable performance across dataset sizes, showing good agreement with the mechanistic model even when trained on only 100 cycles. Nevertheless, in the case of 50 cycle samples, a smaller and more erratic elution peak can be observed, suggesting that the model lacks sufficient data to accurately capture the peak shape.

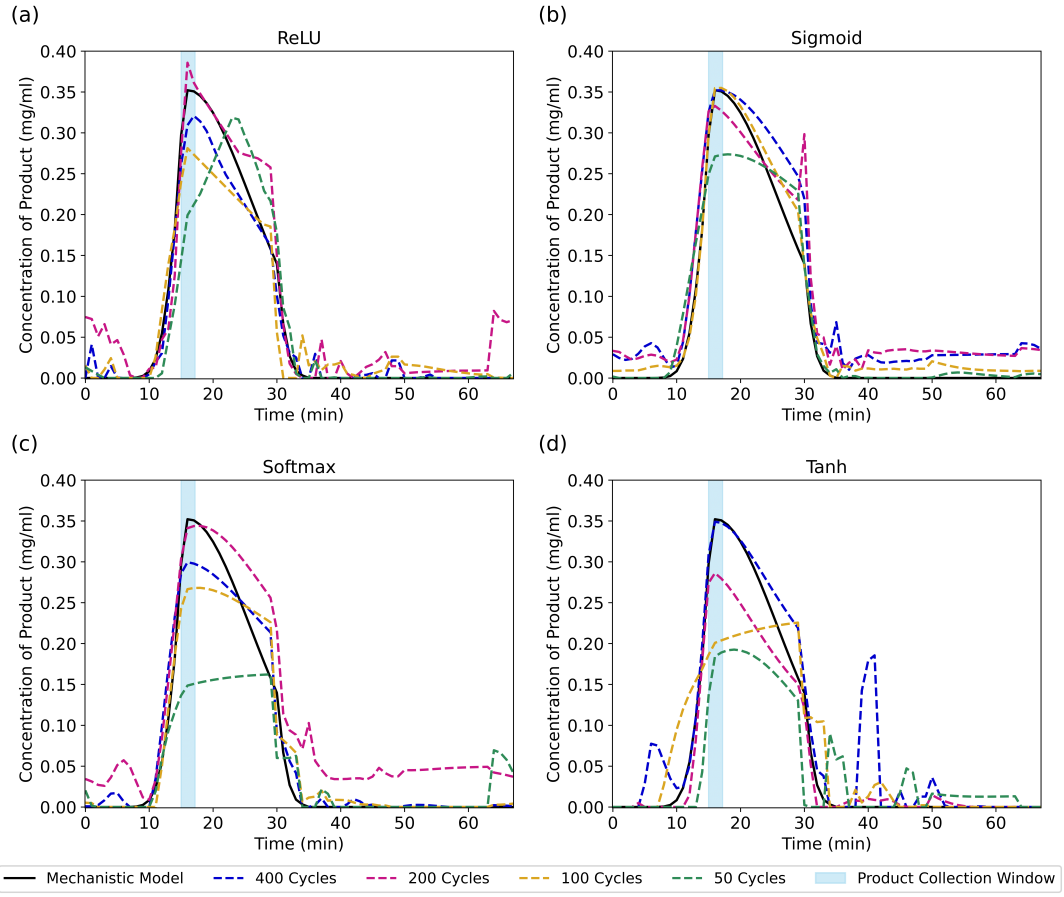


Figure 4.3: Comparison of the outlet product concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

Similarly, the models trained using the Sigmoid activation function manage to capture the height and shape of the elution profile of the product. More specifically, even the 50-cycle model captures the general peak shape and timing with reasonable fidelity. As more training data is introduced, prediction accuracy improves gradually, with the 400-cycle model providing a very accurate representation of the concentration profile. This smooth behaviour is attributed to the saturating nature of the Sigmoid function, which reduces oscillations but can

also dampen fine details when insufficiently trained (Glorot and Bengio, 2010). This results in slightly broader and less pronounced peaks, as the model tends to smooth out steep concentration gradients that define the elution front.

In contrast, the models employing the Softmax activation function display large underpredictions in the elution peak, especially when trained on smaller datasets. Notably, both the 50- and 100-cycle models fail to capture the correct peak height and shape. Only at 400 cycles does the Softmax-based model approach accurate peak and profile characterisation. This underperformance likely stems from the normalisation effects imposed by the Softmax function, which, although useful in classification tasks, may restrict the model's ability to accurately represent high-intensity, narrow peaks, where large variability exists (De Brébisson and Vincent, 2015).

The Tanh-activated models exhibit the largest deviation from the mechanistic model at low dataset sizes. The 50-cycle model produces delayed, fragmented, and oscillatory profiles that do not reflect the expected elution behaviour. This behaviour persist in the 100-cycle model, albeit with reduced severity. At 200 and 400 cycles, the models begin to recover the peak structure, but still underperform compared to the respective ReLU and Sigmoid activated models. This poorer performance may be attributed to the saturation behaviour of the Tanh function at input extremes, combined with the limited representational capacity under data-sparse conditions (LeCun et al., 2012). These observations suggest that model performance cannot be attributed solely to dataset size but is also influenced by the interaction between network architecture, activation dynamics, and the statistical properties of the data.

4.3.1.2 Weak impurities

To further assess the models' ability to generalise across all mixture components, predictions for the weak impurities are now considered. While the product remains the primary target of the separation, evaluating model performance across all species is essential to ensure robustness and interpretability, particularly in identifying potential overlap that may compromise the product purity.

As illustrated in Figure 4.4, the weak impurities are predicted with high accuracy across all activation functions and dataset sizes. The elution profile of this species is characterised by a narrow, early-eluting peak that is well separated

from the product and the strong impurities. This is due to the physicochemical nature of the weak impurities, which demonstrate weaker interactions with the chromatographic column and are therefore not retained in the column to the same extent as the rest of the separation species.

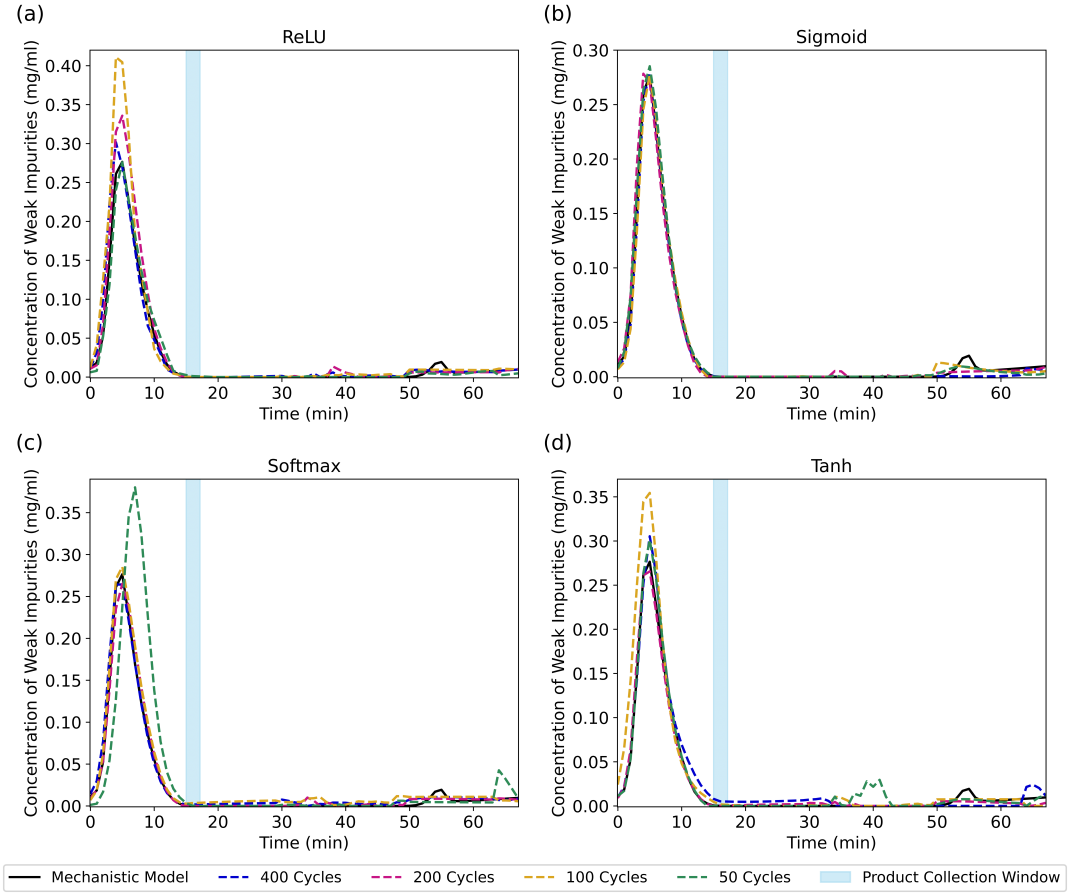


Figure 4.4: Comparison of the outlet weak impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

This spatial and temporal isolation likely contributes to the models' strong predictive performance, as the weak impurities exhibit minimal coupling with the behaviour of the other components. Even in models trained with limited data, such as 50 and 100 cycles, the position and shape of the weak impurities peak are well captured, while some overpredictions of the peak and baseline deviations evident. Notably, in the case of the Sigmoid function, all models predict nearly indistinguishable elution profiles, regardless of data availability during training. In contrast to the more challenging regions of the chromatogram,

where the product and strong impurities are eluted in close proximity, the relative independence of the weak impurities during separation allows the models to learn their dynamics more easily and with less interference from other profiles.

4.3.1.3 Strong impurities

The predictions for the strong impurities, presented in Figure 4.5, show greater variability and sensitivity to model configuration than was observed for the weak impurities. Strong impurities are eluted closer in time to the product and exhibit a broader profile, with significant overlap with the product elution window, due to their strong interactions with the chromatographic column and the product itself. These non-linear interactions increase the difficulty of predicting the different components through the shared layers of the MIMO ANNs.

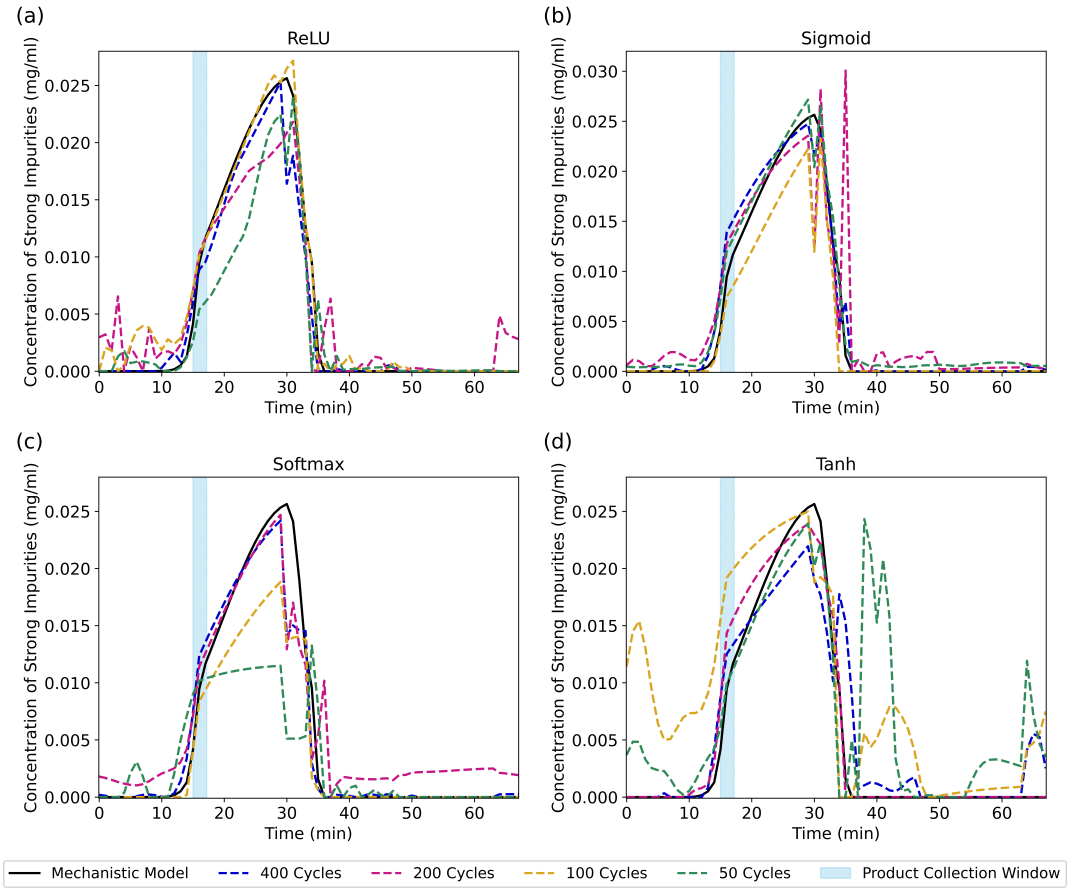


Figure 4.5: Comparison of the outlet strong impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/multiple-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

Models trained with the Softmax and Tanh activation functions on smaller datasets (50 and 100 cycles) exhibit broader, underpredicting peaks and delayed elution. The 50-cycle Tanh model, in particular, shows pronounced oscillations and multiple peaks beyond 35 minutes, which are not present in the mechanistic model, since the column is conducting recycling tasks at that stage and is not eluting any of the components. This suggests reduced stability in the model's internal representation of the species. The instability likely arises from activation function saturation and limited training data, which amplify small numerical fluctuations during prediction. As such, the model appears sensitive to minor variations in the input features, leading to artificial oscillations in regions where the true elution signal should remain flat.

Under data-rich scenarios, the predictions improve across all activation functions, though convergence is not uniform. The 400-cycle ReLU- and Sigmoid-based models capture accurately both peak height and shape. However, even in this case, minor deviations arise, such as the appearance of a secondary peak in the 200-cycle Sigmoid model, thus highlighting persistent challenges in accurately modelling areas of low concentration. This emphasises that model performance for the strong impurities is not only data-dependent but also heavily influenced by component interactions. The overlapping behaviour with the product introduces additional complexity into the model training, which the MIMO structure, with its shared internal representation, may not be suitable to capture under limited data availability.

4.3.1.4 Model performance

To evaluate the interpolation capabilities of the models for all three species of the separation mixture, ten additional datasets are generated by randomly sampling new input values within their validation range. Table 4.3 presents the four performance criteria used in the analysis: (i) the models' accuracy at nominal operating conditions throughout the cycle, (ii) its predictive accuracy within the product collection window, (iii) interpolation accuracy on unseen datasets, and (iv) the computational cost of simulating a single process cycle at CSS.

Table 4.3: Performance of the investigated multiple-input/multiple-output artificial neural networks.

Activation function	Training sample size	Nominal accuracy (%)	Product collection window accuracy (%)	Interpolation accuracy (%)	Single cycle simulation time under CSS (s)
ReLU	50 Cycles	94.5	98.2	92.0	0.9
	100 Cycles	96.1	99.8	94.9	0.8
	200 Cycles	96.5	99.6	94.9	0.9
	400 Cycles	99.0	99.6	94.5	1.0
Sigmoid	50 Cycles	98.7	99.5	98.5	0.9
	100 Cycles	97.6	99.6	98.6	0.9
	200 Cycles	93.7	99.5	96.2	0.8
	400 Cycles	94.6	99.8	97.0	0.9
Softmax	50 Cycles	98.8	97.5	93.9	0.9
	100 Cycles	99.5	98.3	99.5	0.8
	200 Cycles	89.4	99.5	94.1	0.9
	400 Cycles	99.6	99.0	99.5	1.0
Tanh	50 Cycles	86.0	94.2	71.3	0.8
	100 Cycles	88.6	96.4	80.7	0.8
	200 Cycles	93.0	94.3	94.0	0.8
	400 Cycles	98.9	99.6	98.8	1.0

Prediction errors are quantified using the mean absolute percentage error ($MAPE$), as calculated in Equation 4.1, a standard metric in regression analysis.

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{c_{\mathrm{mech}_{i,s,j}} - c_{\mathrm{data}_{i,s,j}}}{c_{\mathrm{mech}_{i,s,j}}}\right| \times 100 \tag{4.1}$$

where, n is the number of data points, $c_{\mathrm{mech}_{i,s,j}}$ is the concentration of component i in stream s during cycle j calculated by the mechanistic process model, and $c_{\mathrm{data}_{i,s,j}}$ is the concentration of component i in stream s during cycle j predicted by the data-driven models. $MAPE$ is calculated by comparing the predicted concentrations of the data-driven models to those obtained from the mechanistic model, providing a consistent measure of accuracy ($100 - MAPE$) across all conditions.

The results in Table 4.3 confirm that the data-driven models generally achieve high predictive accuracy, particularly during the product collection window. The higher accuracy observed within the product collection window compared to the broader elution window primarily reflects how the data-driven models are trained to represent concentration trajectories with a wide dynamic range. During the elution and collection phases, outlet concentrations reach comparatively large values, whereas outside these windows they remain close to zero for most of

the cycle. As a result, prediction errors in high-concentration regions generate larger gradients during optimisation, providing a stronger learning signal and leading the models to represent elution peaks more accurately than near-zero regions. Consequently, model fidelity is highest in regions where concentrations are significant, while deviations in extended low-concentration regions are learned less precisely due to their limited contribution to the loss function during training. This explains the consistently improved agreement observed during the product collection window.

Within this context, differences in model architecture and activation function further influence predictive performance.Notably, the ReLU-based architecture demonstrates consistent performance across all dataset sizes, suggesting that this activation function offers robustness to data sparsity. This can be attributed to the piecewise linear nature of the ReLU function, which promotes sparse and efficient representations and aligns well with the locally affine characteristics of chromatographic elution profiles (Montúfar et al., 2014).

In contrast, the Tanh-based models show a stronger dependency on dataset size. Performance deteriorates significantly as the number of training cycles decreases, particularly in the 50-cycle case. This is likely due to the Tanh function's tendency to saturate at the extremes of its domain, , where gradients approach zero and the network's ability to update its weights is suppressed (Glorot and Bengio, 2010). This can hinder the function's ability to capture sharp transitions in the elution profile when training data are sparse. The non-linearity of the Tanh function requires a denser sampling of the process dynamics to fully characterise input–output relationships. As such, while the function may offer theoretical advantages in modelling non-linear behaviours, its practical effectiveness is constrained by the available training data.

All surrogate models demonstrate good interpolation capabilities when tested against the randomly generated input sets, with the exception of the Tanh-based model trained on 50 cycles. To improve the performance of models trained on this limited dataset, a 5-fold cross-validation approach was applied. However, this did not yield significant improvements in predictive accuracy and increased the computational cost of training by a factor of 24. This suggests that, in this context, cross-validation may not be sufficient to compensate for the lack of data and highlights the limitations of training deep models with limited observations.

The significance of computational efficiency becomes particularly apparent

when considering that the surrogate models predict directly at CSS, bypassing the repeated simulations required by the mechanistic model. While the mechanistic model requires approximately 89 seconds to simulate until CSS is reached (Papathanasiou et al., 2016), the data-driven models are capable of producing predictions in under one second. This corresponds to a reduction in simulation time of 98.9%, representing a substantial computational gain. Such efficiency makes the data-driven models highly attractive for applications requiring iterative simulations, such as online process optimisation and control.

4.3.2 The multiple-input/single-output approach

To evaluate whether independent modelling of the output concentrations mitigates the limitations observed in the MIMO architecture, a MISO strategy is investigated. In this configuration, four separate ANNs are trained, each tasked with predicting a single outlet concentration from the same input set. This decoupling of output prediction aims to eliminate error propagation observed between the predictions of the different species and allow each network to specialise in the representation of a single species' dynamics.

4.3.2.1 Product

The prediction of the product elution profile remains the central focus, as it directly informs both the purity and yield KPIs. Figure 4.6 shows the outlet concentration predictions for the product using the MISO models across four activation functions and decreasing training dataset sizes. In contrast to the MIMO counterparts, the MISO models display consistently lower deviations from the mechanistic model, particularly under limited data availability.

This improvement can be attributed to the decoupling of output predictions. Unlike in the MIMO architecture, where prediction errors in overlapping species could propagate through the shared network layers, the MISO approach enables each model to specialise in learning the dynamics of a single species. This highlights the advantage of learning component-specific representations in isolation, particularly for the product, which shows complex behaviour due to its interactions with both impurities and the solid phase of the columns.

In cases of limited data availability, models trained on 50 cycles using the Sigmoid and Tanh functions still manage to predict the elution window with

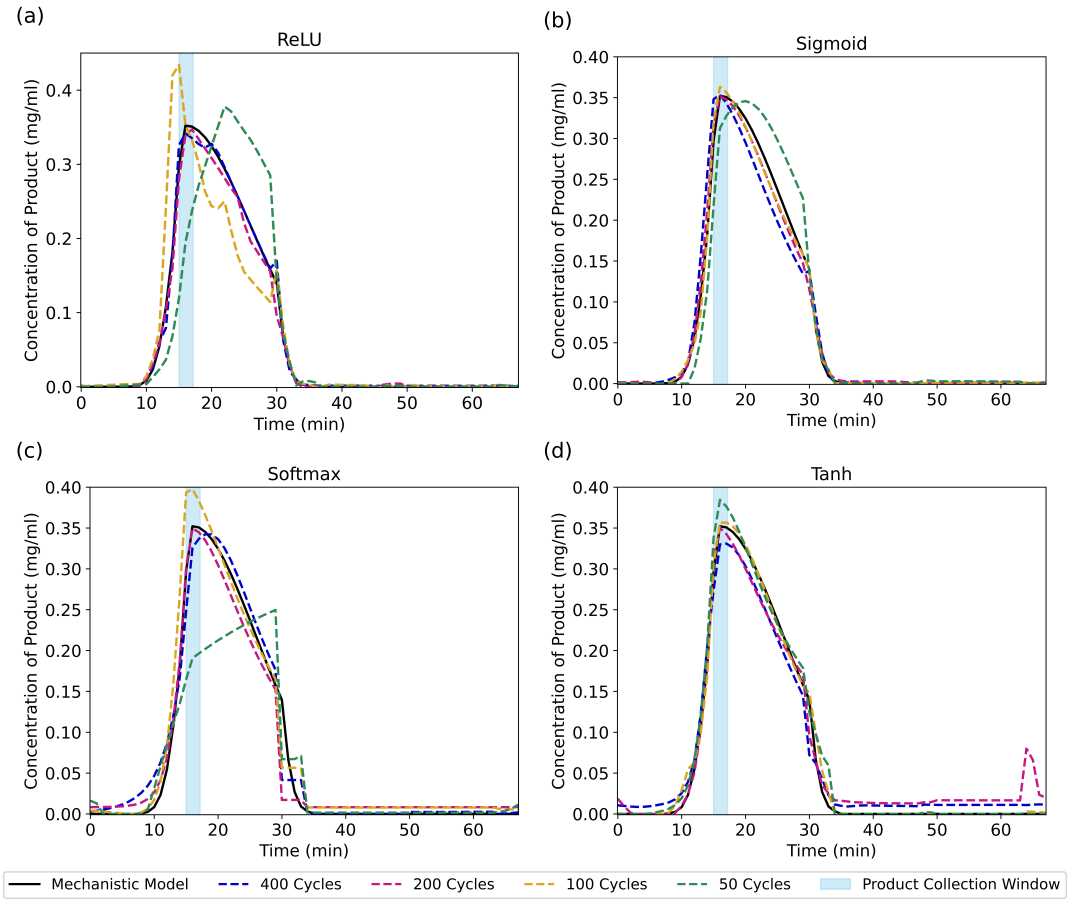


Figure 4.6: Comparison of the outlet product concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

reasonable accuracy, however underestimations and shifts in peak intensity are observed in the cases of the ReLU and Softmax activation functions. In contrast, models trained on 200 or more cycles consistently yield profiles nearly indistinguishable from those of the mechanistic model.

Focusing further on the performance of the ReLU-based model trained on 100 cycle samples, it can be observed that a large deviation occurs during the product collection window. Despite accurate predictions across the rest of the profile, this overestimation of the peak could result in overpredicted KPIs, potentially leading to the collection of an off-spec product. This highlights the need to assess the accuracy of the predicted elution profiles throughout the process cycle, but more specifically during the product collection window, as high accuracy across the full trajectory does not necessarily indicate a reliable model.

From a process perspective, such level of robustness is particularly valuable. In the twin-column MCSGP system (Figure 3.1), the product is collected during well-defined temporal windows under batch-mode operation (B1 phase), and any mischaracterisation of its elution profile can result in the collection of an impure product, which does not meet the required specifications, as well as loss of material. The accuracy of the MISO models within the product elution region therefore provides high confidence in their potential for integration into downstream optimisation or control tasks.

4.3.2.2 Weak impurities

The prediction of weak impurities using the MISO configuration, as shown in Figure 4.7, demonstrates consistently high accuracy across all activation functions and training dataset sizes. As previously observed in the MIMO-based models, the early elution and narrow peak of this separation species make it comparatively easier to model, as its retention behaviour is less affected by inter-component interactions or non-linear process transitions.

Compared to the MIMO models (Figure 4.4), the MISO configurations exhibit improved stability in the baseline region during the second half of the process cycle, during which the column is executing recycling tasks (3.1), particularly when employing the ReLU and Sigmoid activation functions. While these differences are minor, they indicate that the removal of shared internal representations can further reduce propagation of errors in the predicted concentration profiles, even for species that are well separated from the rest of the elution mixture.

In terms of activation function performance, no significant differences are observed, indicating that the prediction of this species is largely insensitive to the non-linearity introduced by the activation mechanism. Likewise, increasing the size of the training dataset does not lead to visible changes in model performance, suggesting that the elution behaviour of the weak impurities can be reliably captured even with limited data. This is consistent with the behaviour of the species, which displays limited interaction with the other mixture components and the stationary phase, and thus presents a simpler learning task for the ANNs.

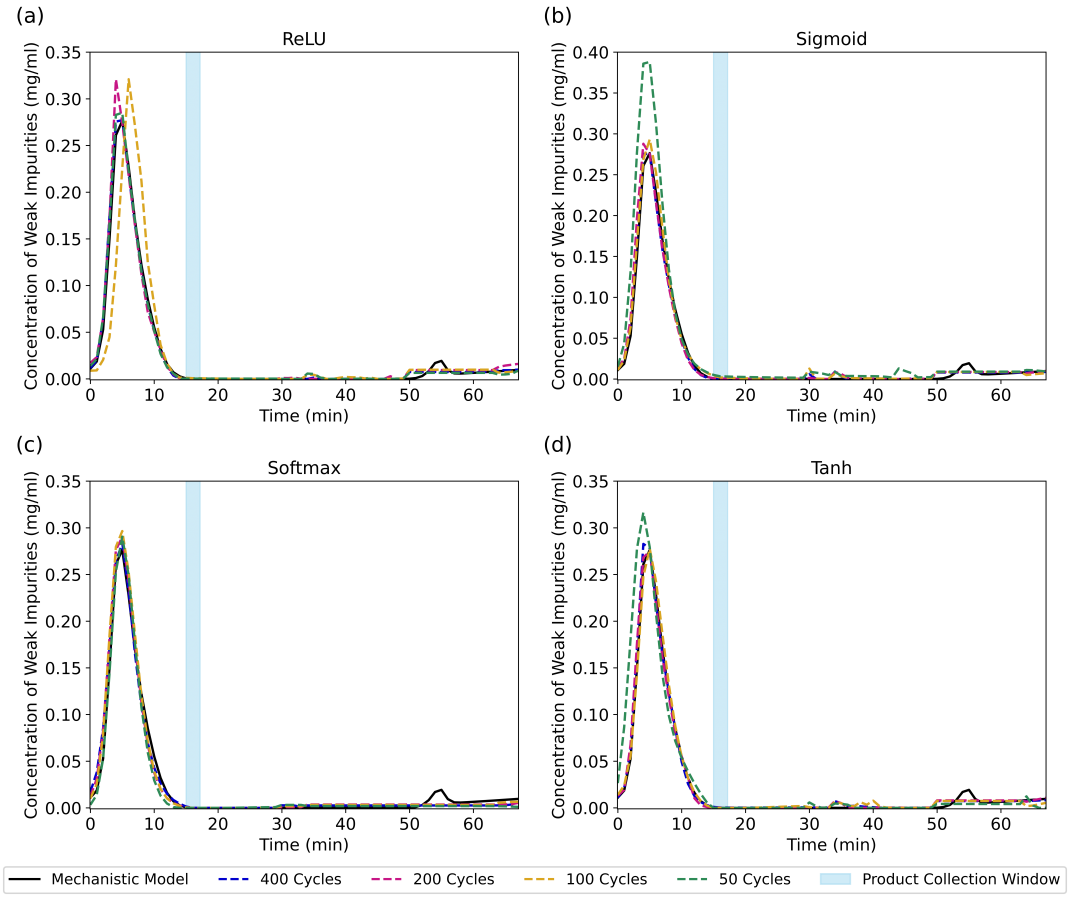


Figure 4.7: Comparison of the outlet weak impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

4.3.2.3 Strong impurities

The strong impurities (Figure 4.8) present a more complex modelling challenge, as they elute later in the cycle and overlap with the product at the end of the B1 phase and the onset of recycling in phase I2 (Figure 3.1). As such, their predictions are sensitive to phase transitions and to co-elution with the product.

Compared to the MIMO models, the MISO models exhibit smoother peak predictions and show a significant reduction in oscillations and secondary peaks. This improvement can be attributed to the elimination of shared weights and nodes in the ANN architecture, which allow for the models to learn the behaviour of the species, without being constrained and influenced by the concurrent learning of the rest of the species.

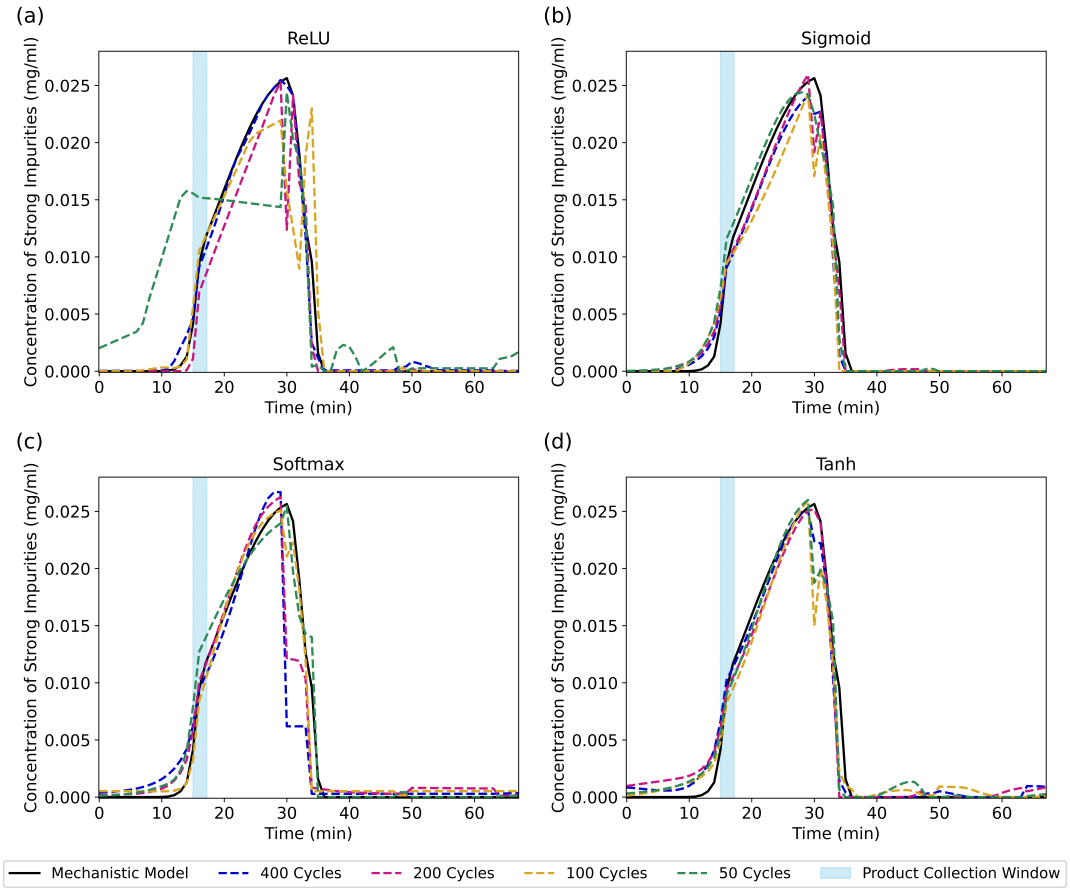


Figure 4.8: Comparison of the outlet strong impurities concentration predictions of the mechanistic and the data-driven models, using multiple-input/single-output artificial neural networks and the (a) ReLU, (b) Sigmoid, (c) Softmax, and (d) Tanh activation functions, for a reducing number of training samples.

This is particularly valuable when training data is limited and sparse, as seen in the 100-cycle MISO models, which produce more stable and accurate predictions compared to even the 200-cycle MIMO models. This implies that decoupling outputs not only improves model fidelity but also reduces data dependency, which is a crucial feature in bioprocess development, where experiments are costly.

4.3.2.4 Model performance

The performance of the MISO models across all evaluation criteria is summarised in Table 4.4. In line with previous observations, prediction accuracy improves with increasing training set size across all activation functions. However, unlike the MIMO case, the effect of data sparsity is more nuanced.

Table 4.4: Performance of the investigated multiple-input/single-output artificial neural networks.

Activation function	Training sample size	Nominal accuracy (%)	Product collection window accuracy (%)	Interpolation accuracy (%)	Single cycle simulation time under CSS (s)
ReLU	50 Cycles	81.2	96.4	61.3	1.1
	100 Cycles	98.9	99.4	59.2	1.0
	200 Cycles	99.8	99.8	95.7	1.0
	400 Cycles	99.6	99.7	98.5	1.1
Sigmoid	50 Cycles	98.4	96.9	98.3	1.1
	100 Cycles	98.1	97.5	98.0	1.1
	200 Cycles	98.3	99.5	97.7	1.1
	400 Cycles	98.0	99.3	98.4	1.0
Softmax	50 Cycles	97.8	99.5	96.9	1.2
	100 Cycles	98.7	99.6	97.1	1.0
	200 Cycles	98.6	99.7	97.7	1.1
	400 Cycles	96.9	99.2	97.5	1.0
Tanh	50 Cycles	96.4	99.6	94.9	1.0
	100 Cycles	97.6	99.3	96.0	1.1
	200 Cycles	96.5	99.5	96.3	1.1
	400 Cycles	96.1	99.5	98.3	1.2

The ReLU-based MISO models, for instance, exhibit a sharp performance drop in interpolation accuracy at lower data volumes, despite showing high accuracy within the product elution window. This suggests that while the model captures the elution peak with sufficient precision to support KPI calculation, it lacks generalisability outside the specific training configurations. A similar pattern is observed for the 100-cycle case, where the model achieves 98% nominal accuracy and 99.4% accuracy within the elution window, yet fails to interpolate effectively across unseen conditions. This highlights the risk of overestimating model reliability based solely on nominal performance metrics.

In contrast, the Sigmoid and Softmax models demonstrate imporved interpolation performance across all dataset sizes, even in low data availability conditions. Both maintain interpolation accuracies above 96% with as few as 50 cycles, indicating better generalisation across the input domain.

Tanh-based models also maintain high elution window accuracy across all sizes, but do not show consistent gains in either nominal or interpolation accuracy as more data are added. This flattening of performance suggests potential saturation in the network's capacity to extract additional value from the enhanced dataset using this activation function.

In terms of computational cost, all MISO models simulate a single cycle

in approximately 1.0–1.2 seconds. Although this represents a marginal increase compared to the MIMO case, the prediction time remains under 2% of the mechanistic model runtime and is therefore well within the bounds required for online integration.

Taken together, the results indicate that the MISO strategy offers a more robust modelling framework than MIMO, particularly when elution window fidelity is prioritised. However, careful consideration must still be given to the interaction between training size and activation function, as strong nominal performance does not always correlate with reliable generalisation. This is especially relevant in process development settings, where untested conditions are frequently encountered. In such cases, the ability of a model to extrapolate beyond the training domain becomes critical, as process datasets are often limited and biased towards nominal operating conditions.

4.4 Conclusions

This chapter presented the development and evaluation of data-driven models for the prediction of outlet concentration profiles in the twin-column MCSGP process, with the objective of enabling fast and accurate simulation under CSS conditions. To this end, MIMO and MISO ANNs were constructed, trained, and assessed against synthetic data generated from an experimentally validated mechanistic model.

Both modelling strategies demonstrated high predictive accuracy within the training domain, with the MISO configuration outperforming MIMO models in terms of elution peak characterisation and robustness to training data scarcity. The comparative analysis of activation functions highlighted ReLU as a suitable choice for sparse, piecewise-linear approximations, whereas functions such as Tanh and Softmax displayed increased sensitivity to data density and output coupling. Despite their structural differences, all networks showed significant computational advantages over the mechanistic model, with speed-ups exceeding 98% in single-cycle CSS simulations.

Nevertheless, several limitations remain. As purely data-driven structures, the models developed rely entirely on the availability and representativeness of the training dataset, and lack embedded knowledge of the underlying process physics. As a result, predictive performance can deteriorate under extrapolative

conditions, while local errors, particularly around the product collection window, can lead to significant deviations in the estimation of the KPIs. Furthermore, the models are constructed under fixed phase durations and do not account for dynamic switching strategies, which limits their application to systems operating under predefined schedules.

To address these issues, the next chapter builds upon the architectures developed here and introduces a hybrid modelling framework that integrates mechanistic knowledge into the model structure. By incorporating physicochemical knowledge and leveraging the mathematical structure of the separation process, the hybrid approach aims to enhance model reliability, extrapolation performance, and physical interpretability, while retaining the computational efficiency required for real-time applications.

Chapter 5

Hybrid modelling

Outputs of this chapter have been disseminated in:

- **F. Michalopoulou** and M. M. Papathanasiou. An approach to hybrid modelling in chromatographic separation processes. *Digital Chemical Engineering*, 14:100215, 2025. (Invited article in the *Emerging Stars* special issue)

5.1 Introduction

The growing demand for efficiency and automation in biomanufacturing has reinforced the role of digitalisation across process development and operation. Digital twins, in particular, have become key enablers, offering dynamic representations of physical systems that can be deployed for simulation, optimisation, and control (Pantelides and Pereira, 2024). In chromatographic separations, these tools are essential to navigate the non-linear dynamics of multicomponent mixtures, and to ensure that strict product specifications are met under industrial constraints (Guiochon et al., 2006; Close et al., 2014). Such requirements demand models that provide high predictive accuracy at a computational cost compatible with repeated or real-time evaluations (Lin et al., 2021).

Mechanistic models remain the reference standard in chromatography, as they provide physically interpretable descriptions of adsorption, transport, and flow phenomena (Kumar and Lenhoff, 2020; Shekhawat and Rathore, 2019). Their formulation through coupled partial differential and algebraic equations (PDAEs), however, requires fine spatial and temporal discretisation, which becomes computationally intensive. In multicolumn configurations operating at cyclic steady state (CSS), the numerical burden can involve solving thousands of equations per iteration, which can constrain their suitability for workflows that require rapid or repeated simulations (Brhane et al., 2019; Rehman et al., 2022).

Data-driven models provide an alternative by learning input–output relationships directly from data (Boukouvala et al., 2011; Solomatine et al., 2008). Methods such as neural networks and Gaussian processes (GPs) have been applied in chromatography to construct surrogates that offer fast predictions and reduced computational burden (Kensert et al., 2022; Ding and Ierapetritou, 2024). Their efficiency makes them attractive for tasks such as monitoring or rapid optimisation, but their dependence on the scope and quality of training data limits generalisability. Without embedded physical constraints, such models extrapolate poorly and provide limited interpretability, which restricts their adoption in biopharmaceutical contexts where robustness and physical consistency must be demonstrated (Anupa et al., 2025; Yang et al., 2023).

Hybrid models have been proposed to address these limitations by embedding selected mechanistic relationships into machine learning frameworks. In doing so, they retain the interpretability and robustness of first-principles models while

leveraging the flexibility of data-driven approaches (Ding et al., 2023; Narayanan et al., 2021a). These models are particularly useful when process knowledge is incomplete or when experimental data are scarce (Agharafeie et al., 2023). They also enable transfer learning, where knowledge obtained from one system can be adapted to another with similar characteristics, such as different column properties or product variants (Kay et al., 2023; Sansana et al., 2024).

Hybrid models in chromatography have been applied to tasks ranging from embedding adsorption relationships to parameter calibration and retention modelling. Neural networks with integrated kinetic, equilibrium, or retention expressions have been proposed to preserve mechanistic structure while introducing predictive flexibility (Nagrath et al., 2004; Pirrung et al., 2017). Other studies have employed hybrid frameworks for calibration and root-cause analysis (Wang et al., 2017a,b), while more recent formulations have accurately represented equilibrium and retention behaviour (Narayanan et al., 2021a; Regufe et al., 2021).

Despite recent advances, the majority of reported hybrid models in chromatography address parameter estimation or the optimisation of scalar key performance indicators (KPIs), such as yield, purity, and productivity. Approaches that extend to dynamic prediction remain comparatively limited, and consistent generalisability across varying feed compositions and broader operating domains has yet to be demonstrated (Agharafeie et al., 2023; Narayanan et al., 2021a). Moreover, although hybrid frameworks can reduce the computational burden of mechanistic models, their deployment in real-time environments still depends on demonstrating sufficient numerical efficiency, stability, and robustness.

In this Chapter, a hybrid modelling framework is developed for the twin-column multicolumn countercurrent solvent gradient purification (MCSGP) process. A mechanistic description of adsorption equilibria is combined with artificial neural networks (ANNs) trained to approximate elution concentrations at CSS, thereby eliminating the need for spatial discretisation while retaining predictive accuracy. The formulation is expressed entirely in terms of experimentally measurable variables, which enhances transparency and facilitates practical deployment. The model is evaluated through validation, interpolation, and extrapolation studies, providing a systematic assessment of its predictive performance and robustness. The results illustrate how hybrid modelling can provide an efficient and reliable representation of non-linear separation dynamics, establishing a foundation for future integration into optimisation and control tasks.

5.2 Methodology

The hybrid model presented in this Chapter is developed following four main steps: (i) system analysis, (ii) model architecture design, (iii) synthetic data generation, and (iv) surrogate architecture identification and training. The methodologies underlying each of these steps are described in the following sections.

5.2.1 System analysis

The starting point for developing the hybrid model is a careful analysis of the underlying dynamics of the twin-column MCSGP system. As discussed in Chapter 3, the process involves the separation of a ternary mixture under cyclic operation, where the non-linear interactions between species and the modifier lead to highly complex behaviour. To design a model that is both accurate and computationally efficient, it is necessary to first identify the main sources of complexity in the governing equations and to determine which mechanistic elements must be retained in reduced form to preserve predictive fidelity.

The full mechanistic process model is formulated as a set of coupled PDAEs, which describe the evolution of liquid- and solid-phase concentrations (Equations 3.1 and 3.2) along the column length. The dynamic balances for the mobile and stationary phases introduce spatial dependence, and the numerical solution of these equations requires spatial discretisation into collocation points. This discretisation greatly increases the number of variables and equations to be solved per iteration, and is the principal contributor to the computational cost of the mechanistic framework.

A closer inspection of the mass balances further reveals the presence of a feedback loop between the governing equations. At each time step (t), the concentration of solute in the mobile phase ($c(z,t)_{i,h}$) is used to compute its equilibrium solid-phase concentration ($q^*\left(c(z,t)_{i,h}\right)$) via the separation isotherm (Equation 3.6). This equilibrium term, in turn, determines the rate of mass transfer into the stationary phase ($\frac{\partial q(z,t)_{i,h}}{\partial t}$), which feeds back into the balance for the liquid phase concentrations at the next time step ($t+1$). Such nested dependencies highlight the strongly coupled nature of the dynamics and the central role of adsorption in shaping the outlet profiles.

The most critical mechanistic feature in this chain of interactions is the com-

petitive adsorption behaviour described by the isotherm. In the mechanistic model, the solid-phase equilibrium concentration of each component is expressed as a function of the concentrations of all species in the mobile phase, modulated by the Henry constants and saturation capacities of the adsorption sites (Equations 3.8-3.12). These parameters themselves are dependent on the instantaneous concentration of the modifier, embedding the effect of the modifier gradient into the adsorption dynamics. As a result, the isotherm captures the essential non-linear competition between product and impurities, and provides the mechanistic basis for the separation achieved in the twin-column MCSGP system.

For these reasons, the isotherm is preserved in the hybrid framework. Retaining this equation ensures that the competition for binding sites and the influence of the modifier are represented explicitly, even as the spatial discretisation of the column is removed. In this way, the hybrid model eliminates the most computationally demanding aspect of the mechanistic formulation while maintaining the fundamental physics of the separation. This balance between simplification and fidelity provides a foundation for the development of a surrogate model that is computationally efficient yet capable of capturing the critical dynamics of the process.

5.2.2 Model architecture

Based on the model and equation analysis, a hybrid modelling structure is defined, in order to reduce the model complexity, while maintaining process knowledge. The hybrid structure therefore retains the equilibrium isotherm (Equation 3.6), while the spatially distributed mass balances (Equations 3.1 and 3.2) are replaced by a data-driven surrogate. In this way, the formulation removes the computational burden associated with discretising and solving coupled PDAEs, but continues to preserve the adsorption physics that determine selectivity and peak overlap.

The input–output structure of the model is defined by extending the formulation used for the data-driven models in Chapter 4. The operating conditions of the process, namely the feed composition, flowrate, and initial modifier concentration, are included as model inputs alongside the instantaneous inlet concentrations of the modifier and separation species at each discrete timestep of one

minute. The one-minute resolution is consistent with the characteristic dynamics of the twin-column MCSGP process and with the choice previously established for the data-driven models, and allows the hybrid formulation to capture transient elution features while maintaining computational efficiency. The surrogate is tasked with predicting directly the outlet concentrations of the modifier, weak impurities, product, and strong impurities at CSS, thereby removing the need for spatial resolution within the column.

Further to these inputs, an additional dynamic feature is included to incorporate physicochemical knowledge of the system (Figure 5.1). As discussed in the previous section, a feedback loop is formed between the liquid-phase concentrations, the equilibrium adsorption, and the update of the mobile phase in the next timestep. To preserve this dependency in reduced form, the equilibrium concentrations of the species at the outlet from the current timestep ($q^*\left(c(z=L,t)_{i,h}\right)$) are supplied as inputs to the ANN, in order to predict the respective outlet liquid phase concentrations of the species at the following timestep ($c(z=L,t+1)_{i,h}$). These values are calculated through the isotherm retained from the mechanistic model, with parameters defined either experimentally or as functions of the modifier concentration. Including these equilibrium terms ensures that the surrogate retains a memory of the adsorption state and captures the sequential nature of elution under gradient conditions without the need to solve the full set of mass balances.

In practice, an important consideration in the design of the hybrid model is that it can operate with different levels of measurement availability. When detailed chromatogram data are available, the equilibrium features can be calculated directly from measured outlet concentrations. When only pooled fraction data or absorbance signals are available, the ANN can calculate the equilibrium inputs using its own predicted outlet concentrations from the previous timestep. This construction ensures that the model is robust to real-world scenarios, where the resolution of available measurements may vary. In particular, UV absorbance at 300 nm, which is routinely recorded in twin-column MCSGP systems, provides sufficient information to support the implementation of the hybrid model in practice.

Since the hybrid surrogate does not include the spatially distributed mass balances, conservation of mass is not enforced by model design. To address this limitation, a cycle-level correction is applied after each cycle prediction. The

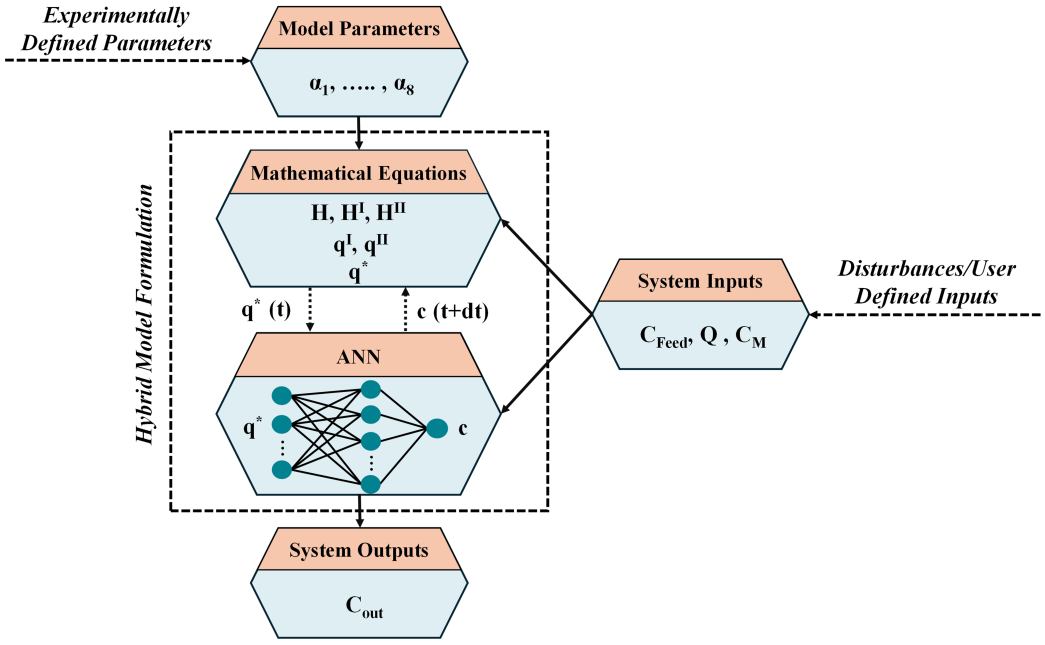


Figure 5.1: Schematic overview of the structure of the hybrid model.

total outlet mass of each species is compared with the inlet mass over the same cycle, and the resulting imbalance is redistributed across the predicted profile in proportion to the local predicted elution concentrations. In this way, the correction follows the dynamics of the profile and preserves peak shape while enforcing mass balance. This adjustment is computationally inexpensive and restores mass balance, as shown in Equation 5.1.

$$\begin{aligned} c(z=L,t)_{i,h}^{\text{corrected}} &= c(z=L,t)_{i,h} \\ +\frac{1}{Q_h}\left(\frac{c(z=L,t)_{i,h}\cdot Q_h}{\sum_{t=0}^{t_{\text{end}}}\left(c(z=L,t)_{i,h}\cdot Q_h\right)}\sum_{t=0}^{t_{\text{end}}}\left(Q_h\left(c(z=0,t)_{i,h}-c(z=L,t)_{i,h}\right)\right)\right) \end{aligned} \tag{5.1}$$

where $c(z=L,t)_{i,h}^{\text{corrected}}$ is the corrected eluted liquid phase concentration of component i in column h. The corrected outlet concentration of each component is obtained by adding an adjustment term to the predicted concentration at the column outlet ($c(z=L,t)_{i,h}$). This adjustment term is proportional to the mass imbalance calculated over the full cycle, defined as the difference between the inlet ($c(z=0,t)_{i,h}$) and outlet ($c(z=L,t)_{i,h}$) concentrations of the component, and scaled by the corresponding flowrate (Q_h). In this way, the imbalance is redistributed dynamically along the elution profile. Larger corrections are applied in regions where the elution signal is higher, while smaller adjustments are introduced in regions of low elution. This preserves the overall shape of the predicted peaks, while ensuring that mass conservation is satisfied at the cycle level.

These elements define the hybrid model architecture. By retaining the isotherm, the formulation preserves the physics of competitive adsorption, while the ANN provides the flexibility to approximate the outlet mapping. Further to that, the inclusion of equilibrium inputs ensures mechanistic memory of adsorption state, and the mass-balance correction restores physical consistency at the cycle level. In this way, the hybrid model eliminates the need for spatial discretisation, while maintaining the accuracy required to describe the process at CSS.

5.2.3 Synthetic data generation

For the training of the hybrid model, synthetic data are generated by varying the nine input variables that describe the process conditions of interest, namely the feed concentrations of all species, the flowrates applied during the process phases, and the initial modifier concentrations in both columns, similarly to the data generation performed in Chapter 4. The specific ranges considered are listed in Table 5.1. These ranges are chosen to remain within the range of validity of the mechanistic model. The variation of the inputs is performed using a quasi-random Sobol sequence, which ensures an even exploration of the multidimensional input space and avoids clustering effects associated with purely random sampling (Sobol' and Shukman, 1993). For the feed concentrations, the variation is restricted to 50% of the full mechanistic validation range. This design allows for broader validation of the resulting model, as it enables interpolation to be evaluated within this training region, while extrapolation can be assessed by testing the trained model on operating points outside it.

Table 5.1: Data generation inputs.

Input	Lower bound	Upper bound
Feed concentrations		
Feed concentration of weak impurities (mg/mL), ($C_{\text{feed}W}$)	0.065	0.075
Feed concentration of product (mg/mL), ($C_{\text{feed}P}$)	0.35	0.45
Feed concentration of strong impurities (mg/mL), ($C_{\text{feed}S}$)	0.035	0.045
Flowrates		
Flowrate during phases B1 and B2 (mL/min), (Q_B)	0.1	1
Flowrate during phases I1 (mL/min), (Q_{I1})	0.1	1
Flowrate during phases I2 (mL/min), (Q_{I2})	0.1	1
Modifier concentrations		
Feed concentration of modifier (mg/mL), ($C_{\text{feed}M}$)	1	3
Initial modifier concentration of Phases I1 for column 1 (mg/mL), (C_{M1})	2	3
Initial modifier concentration of Phases I1 for column 2 (mg/mL), (C_{M2})	0.8	1.2

For each operating point sampled, the mechanistic model is simulated until CSS is achieved, and outlet concentration profiles of the modifier, weak impurities, product, and strong impurities are collected at one-minute intervals. The corresponding inlet concentrations and equilibrium adsorption terms are also extracted, providing the complete set of input–output pairs required for training (Figure 5.2). This approach embeds the system dynamics into the dataset implicitly, as the surrogate is informed of variable inlet profiles, which indirectly encode time dependence, without requiring explicit temporal information. In this way, the model is not constrained to assign elution peaks to fixed time points, but can instead reproduce the shifting of elution windows under different operating conditions.

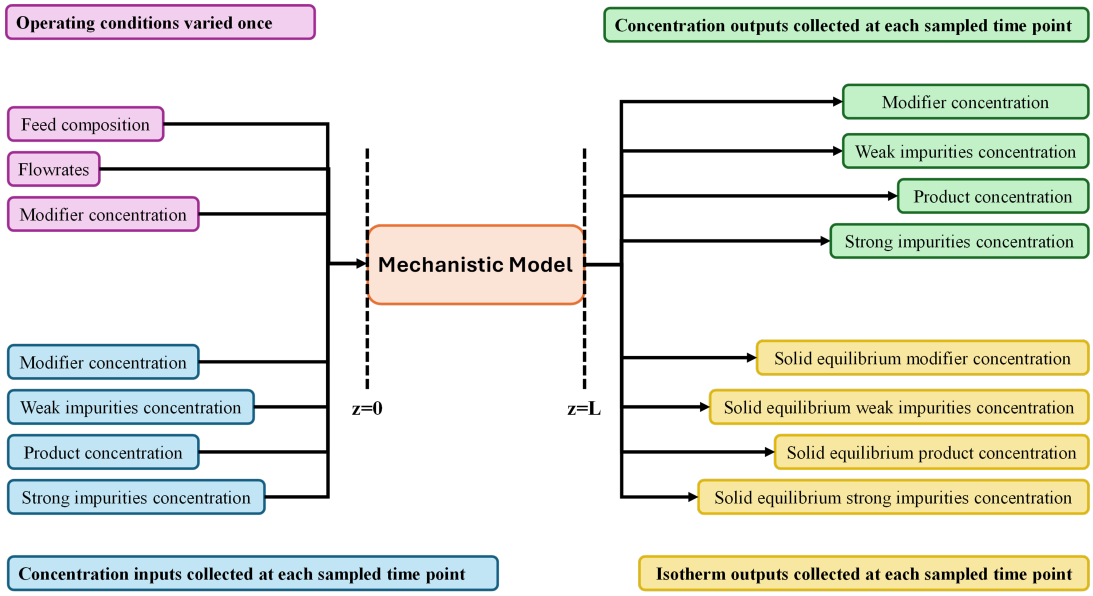


Figure 5.2: Schematic representation of the input–output structure for the synthetic data generation.

The elimination of explicit temporal encoding means that phenomena such as severe column overloading or strong mixing effects are not directly represented (Dai et al., 2009; Gritti and Guiochon, 2012). However, these effects lie outside the validation range of the mechanistic model and are therefore not relevant for the present application. Within the operating domain of interest, where the hybrid model is designed to operate, the surrogate remains accurate under CSS conditions. This methodology aligns with the principles of feature engineering, where rather than leaving the network to infer all relationships from raw inputs, domain knowledge is explicitly incorporated to guide learning and improve

predictive performance (Amos and Kobayashi, 2019).

In total, 400 different operating points are simulated, resulting in a dataset of 27,200 input–output pairs. These 400 cases are divided into four subsets, with 100 cases varying only the feed concentrations, 100 cases varying only the flowrates, 100 cases varying only the modifier concentrations, and 100 cases varying all three classes of variables together. This structure reflects the fact that in chromatographic systems, feed composition, flowrate, and modifier concentration each influence elution behaviour individually, while their combined variation produces strong interactions. Including both types of variation ensures that the hybrid model is trained to capture the singular and interactive effects that govern separation performance. The dataset is divided into training and testing sets. The training set is used to fit the network parameters, while the testing set to guide hyperparameter selection and prevent overfitting. All variables are normalised prior to training using consistent scaling factors across all subsets, to improve numerical stability.

The modifier gradient and step durations are kept constant, and start-up and shut-down cycles are excluded. This decision is supported by sensitivity studies reported in the literature, which showed that flowrates and modifier levels are the most critical parameters affecting separation outcomes (Papathanasiou et al., 2019b). This observation is also consistent with the system analysis in Chapter 3, where variations in flowrate, feed composition, and modifier concentration were found to drive the largest changes in elution behaviour. On this basis, the training dataset is restricted to the most influential operating conditions, avoiding unnecessary expansion of the input space while preserving predictive fidelity.

By generating the dataset in this structured way, the hybrid model achieves three key objectives. It relies solely on measurable or computable inputs, it maintains computational efficiency during training, and it adapts flexibly to shifting elution behaviour without requiring explicit time encoding. These properties ensure that the hybrid model remains both accurate within the validated operating space and robust when applied to less familiar scenarios outside it.

5.2.4 Surrogate architecture and training

The data-driven component of the hybrid model is implemented as a multiple-input/multiple-output (MIMO) ANN. A single ANN is used to approximate the four monitored outputs, namely the outlet concentrations of the modifier, weak impurities, product, and strong impurities, based on the defined set of inputs.

The resulting system corresponds to a 17–4 mapping, where the 17 inputs consist of the nine operating conditions of the system (feed composition, flowrates, and initial modifier concentrations), the four inlet concentrations of the modifier and solutes, and the four calculated equilibrium solid-phase concentrations at the outlet. This choice reflects the coupled nature of chromatographic separation, where deviations in one component profile directly influence the others. By enforcing joint prediction, the surrogate is able to capture peak overlap and relative migration, which are central to the yield–purity trade-off. Unlike recurrent neural networks (RNNs), the hybrid model does not require explicit temporal encoding. The combination of the evolving inlet profiles, the operating conditions, and the equilibrium features is sufficient to embed the required dynamic behaviour, while avoiding the complexity and stability issues that often arise when recurrent network architectures are used. (Figure 5.3).

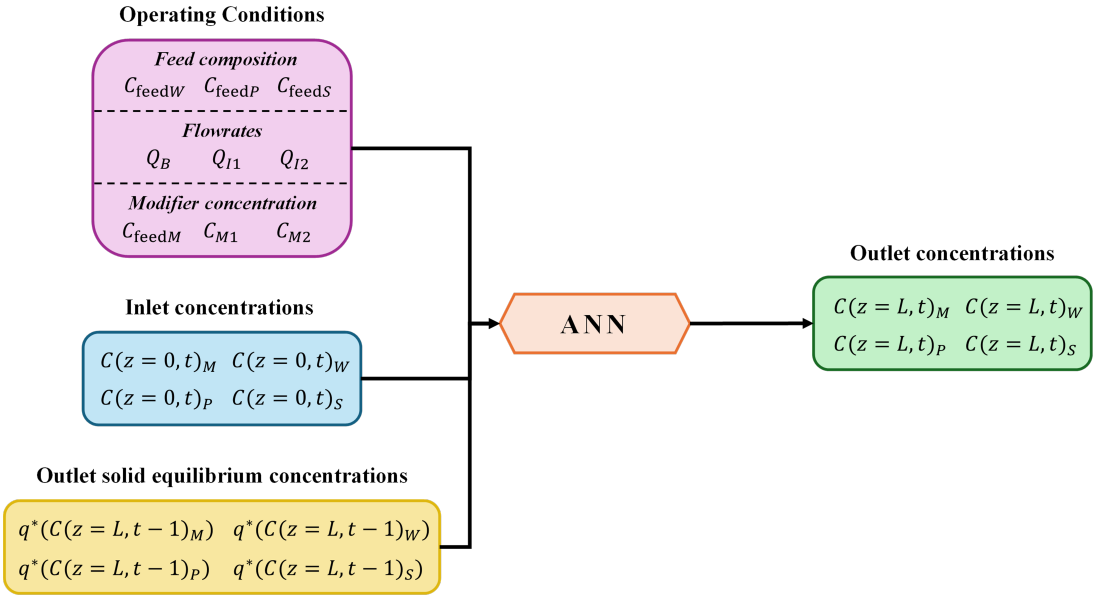


Figure 5.3: Schematic representation of the input–output structure of the ANN used in the hybrid model.

Hyperparameter tuning is performed using Bayesian optimisation (BO). The

hyperparameters considered are the number of hidden layers, the number of neurons per hidden layer, and the learning rate. These variables are selected because they directly determine the expressiveness of the network and the rate of convergence during training. The corresponding search bounds are shown in Table 5.2, and are based on ranges reported effective in previous development of data-driven models for the twin-column MCSGP process (Chapter 4). BO is preferred over grid search or random search because it balances exploration and exploitation in the hyperparameter space, allowing good configurations to be identified with fewer training runs. This efficiency is essential given the high dimensionality of the space and the computational cost of training neural networks to convergence. The objective function for the optimisation is the mean squared error (MSE) between predicted and reference concentrations, which directly quantifies the deviation of the surrogate from the mechanistic model.

Table 5.2: Hyperparameter bounds for the Bayesian optimisation.

Hyperparameters	**Lower bound**	**Upper bound**
Hidden layers	1	5
Neurons per hidden layer	16	128
Learning rate	10^{-5}	0.1

The ANN is constructed with a uniform architecture, such that all hidden layers contain the same number of neurons. This design choice reduces the dimensionality of the hyperparameter space while retaining sufficient flexibility to approximate the highly non-linear dynamics of the chromatographic process. ReLU activation functions are employed in all hidden layers, as they are widely used for their stability and efficiency in deep learning applications, and proved to be most effective in the development of data-driven models for the system (Chapter 4). Weight initialisation is performed using Kaiming Uniform initialisation, with biases set to zero, which has been shown to improve the convergence behaviour of ReLU-activated networks by maintaining a consistent variance of activations across layers (Narkhede et al., 2022; Wong et al., 2024). To control computational demand during the BO procedure, each candidate architecture is trained with a single weight initialisation.

Training is carried out using the Adam optimiser, which provides adaptive learning rates and has been demonstrated to perform robustly across a wide range

of deep learning applications. Each network is trained for a maximum of 40,000 epochs, with early stopping applied if the validation loss does not improve for 500 consecutive epochs. This prevents overfitting and reduces unnecessary computation when a given architecture ceases to improve. The training data consist of 80% of the synthetic dataset, with 20% reserved for testing the hyperparameters during training. Hyperparameter tuning and early stopping are therefore performed iteratively until the optimal configuration is found. All variables are normalised consistently across these subsets to improve numerical stability.

5.3 Results & discussion

5.3.1 Hybrid model validation

A hybrid model is developed based on the presented methodology, retaining the equilibrium isotherm and embedding an ANN to approximate the outlet concentrations at CSS. The BO-identified optimal configuration of the network comprises 3 hidden layers, 98 neurons per layer, and a learning rate of $8.04 \cdot 10^{-3}$.

The predictive accuracy of the hybrid model is first examined under nominal operating conditions, and is validated against the mechanistic model of the system. The predicted profiles match the mechanistic model closely across all components (Figure 5.4). The product peak is well captured, particularly in the region of the product collection window, which is the most critical part of the cycle since product purity and process yield are evaluated directly from this fraction. In this region, the hybrid model reproduces both the height and shape of the elution peak, ensuring that the amount of product recovered is correctly represented. The description of the impurities around the product window is also accurate. The tail of the weak impurities is slightly underpredicted in intensity, but the decay towards the baseline aligns with the mechanistic model, meaning that their contribution within the product collection window is correctly represented. On the opposite side, the rise of the strong impurities is predicted with the right timing and progression, preserving the onset of overlap with the product. By maintaining these boundaries, the hybrid model reflects the same yield–purity trade-off that governs the separation.

The remaining parts of the cycle are also reproduced with high fidelity. The main peak of the weak impurities that elutes before the product is underpredicted

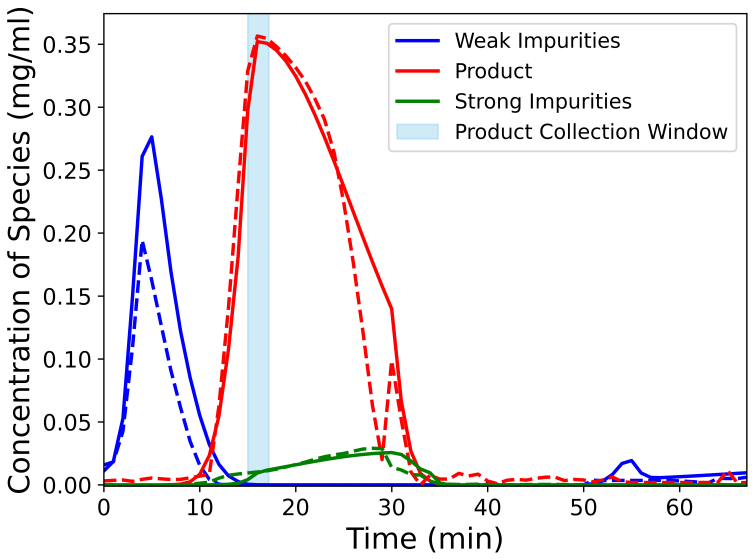


Figure 5.4: Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under nominal operating conditions.

in height and narrower than in the mechanistic model, but its overall position and return to baseline are consistent. This ensures that the product window is not affected by misplaced weak-impurity residual overlap. The strong impurities peak following the product is predicted with the correct magnitude and decay, maintaining the overlap on the trailing side of the product window. Although these regions are less critical for the direct calculation of KPIs, they remain important for preserving the overall material balance of the system. The correct sequencing of weak impurities, product, and strong impurities across the cycle confirms that the hybrid model retains the timing of elution events, which is essential for representing the cyclic operation of the process.

The agreement observed in the nominal profiles is confirmed by the performance metrics reported in Table 5.3, where both the root mean square error (RMSE) and the mean absolute percentage error (MAPE) are used to quantify deviations from the mechanistic model. RMSE captures the absolute magnitude of deviations, while MAPE reflects relative accuracy ($100 - MAPE$) by expressing errors as a percentage of the reference values (Equation 4.1). Across the full cycle, the RMSE remains low for all three species, with the largest deviation corresponding to the weak impurities, consistent with the underprediction of their main peak. The product and strong impurities exhibit smaller errors, reflecting the close match in both peak magnitude and timing. When the analysis is re-

stricted to the product collection window, the errors decrease further. For the weak impurities, the residual tail within the collection window is predicted at the correct low level, ensuring that their contribution to the collected fraction is represented consistently with the mechanistic model. For the product, the RMSE remains small, confirming that the peak height and shape are reproduced with the precision required for accurate yield estimation. The strong impurities also show low error values within the window, indicating that their overlap with the product is represented correctly in both magnitude and timing.

Table 5.3: Prediction performance metrics under nominal operating conditions.

Evaluation Window	**Metric**	**Weak Impurities**	**Product**	**Strong Impurities**
Process cycle at CSS	RMSE (mg/mL)	0.026	0.027	0.0033
	Accuracy (%)	91	97	94
Product collection window	RMSE (mg/mL)	0.00024	0.016	0.0027
	Accuracy (%)	99	99	98

The distribution of errors between the full cycle and the product collection window highlights the role of the isotherm retained in the hybrid formulation. While the network underestimates the main peak of the weak impurities, it still aligns the decay correctly, which limits the impact on the KPI-relevant region. The product and strong impurities predictions remain stable across both scales of evaluation, showing that the features most critical to purity and yield are preserved. This confirms that the hybrid model reproduces the mechanistic profile with high accuracy at the cycle level, while ensuring that the calculated KPIs are not distorted by local prediction errors.

Finally, the computational efficiency of the hybrid formulation is benchmarked against the mechanistic model. The replacement of the spatial balances with the ANN reduced the cost of predicting a CSS cycle by 97%, while retaining accuracy in both profile reconstruction and KPIs calculation. This difference is significant, as the mechanistic model requires the repeated solution of thousands of PDAEs per iteration, and numerous cycles until CSS is achieved, whereas the hybrid formulation provides direct predictions at CSS. The reduction in runtime allows the hybrid model to be evaluated repeatedly within a feasible timescale, making it suitable for applications where repeated evaluations and rapid assessment of process performance is required.

5.3.2 Interpolation

The interpolation performance of the hybrid model is assessed using four test scenarios within the training domain, where the feed composition is varied by ±10% and the flowrates by ±20%, compared to their respective nominal values. These perturbations probe the ability of the model to adjust to shifts in peak timing and intensity while remaining within the range of operating conditions seen during training. The resulting outlet concentration profiles are compared against the mechanistic model across one full cycle at CSS.

The first set of tests considers variations in the feed concentrations of the separation species (Table 5.1). Since feed concentration is a measured disturbance in practice, the model must be able to adapt to such variations if it is to provide reliable predictions across operating cycles. At low concentrations of feed (Figure 5.5a), the amplitudes of all species decrease relative to the nominal case, while their elution timing remains unchanged. The hybrid model follows this scaling closely. The product peak is reproduced with the correct height and width, ensuring that the calculated recovery remains consistent. Similarly, the weak impurities are accurately predicted, with the hybrid model managing to capture both the height and timing of the elution peak. The main deviation arises in the strong impurities, which are underpredicted in intensity, though their timing remains correct. This underestimation, however, does not extend into the product collection window and therefore has limited impact on the process KPIs.

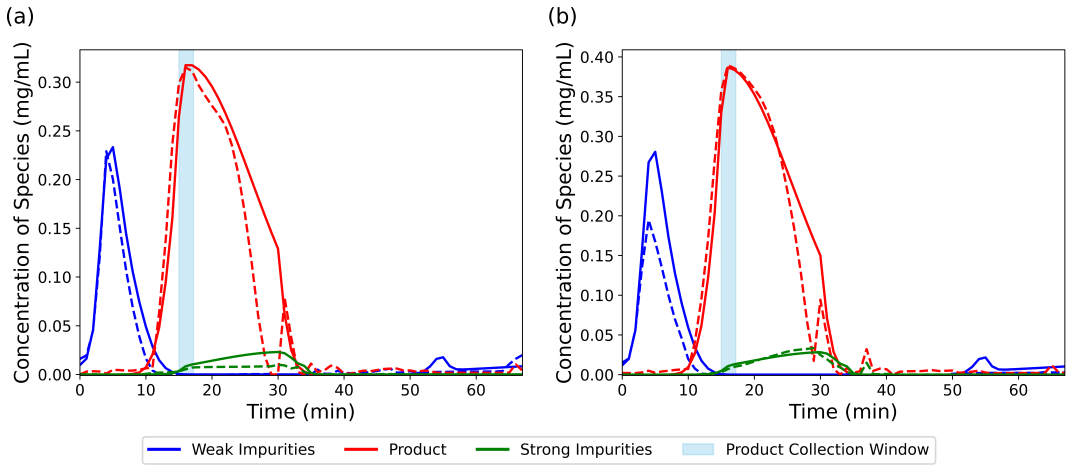


Figure 5.5: Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under interpolation conditions of feed variation: (a) −10% feed and (b) +10% feed.

At high concentrations of feed (b), the amplitudes increase across the cycle. The hybrid model captures this behaviour, reproducing the higher product peak without distorting its width or position within the collection window. However, a small additional peak appears on the trailing edge of the product in the hybrid prediction, which is not present in the mechanistic model. While this artefact does not affect the main peak, it shows that the network introduces a local deviation under increased load. Despite this, the principal product peak and its overlap with the strong impurities are captured with sufficient accuracy to preserve the KPIs calculation. The weak impurities remain well aligned in timing and shape, though slightly underpredicted in height, while the strong impurities follow the mechanistic model more closely than in the low-feed case.

The second set of tests considers variations in the flowrates applied at each process phase (Table 5.1). Unlike feed composition, which acts as a measured disturbance and primarily affects the height of the elution peaks, changes in flowrate directly alter the dynamics of peak migration and shape. At low flowrate values (Figure 5.6a), the peaks broaden and elute later in time, as expected. The hybrid model captures this behaviour across all components. The product peak is predicted with the correct delay and width, ensuring that the collected fraction is represented consistently, as is the case for the strong impurities as well. In contrast, the weak impurities are slightly underpredicted in height but remain aligned in timing, while their decay into the baseline is matched to the mechanistic model.

At high flowrate values (Figure 5.6b), the peaks shift earlier and become narrower. The hybrid model reproduces this shift, capturing the advance of the product peak and its compressed width within the product collection window. The weak impurities follow the mechanistic model closely in both timing and amplitude, preserving the overlap with the product elution window. The strong impurities are predicted with the correct onset but with a distorted profile, being slightly overestimated at the start of elution and underestimated thereafter. In addition, the hybrid introduces a small overprediction in the tail of the elution curve. This overprediction however does not affect the KPIs calculation, as it takes place after the product collection window. These results show that the hybrid model adapts to flowrate variations by reproducing the timing and shape changes that arise from altered transport dynamics, maintaining fidelity where it matters most for KPIs calculation.

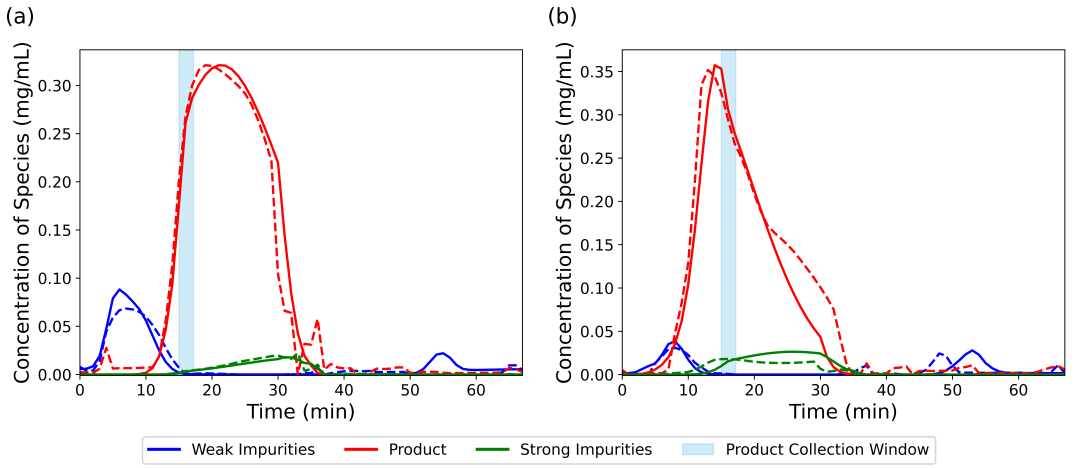


Figure 5.6: Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under interpolation conditions of flowrate variation: (a) −20% flowrate and (b) +20% flowrate.

The agreement observed in the interpolation profiles is reflected in the accuracy metrics reported in Table 5.4. Across all scenarios, the RMSE values remain low, with the largest deviations linked to the weak impurities peak outside the product collection window. Within the product collection window, both RMSE and percentage accuracy confirm that the hybrid model reproduces the mechanistic profiles with high fidelity. For the feed variation cases, the scaling of peak amplitudes is captured without distorting the local overlaps, so the calculated KPIs remain consistent with the mechanistic model. For the flowrate variation cases, the shifts in elution time and peak width are followed accurately, and the small deviations observed in the weak and strong impurities profiles do not translate into KPI-errors. Accuracy values exceed 93% across the cycle and reach 98% in the collection window, showing that the hybrid model maintains reliable predictive performance under the operating variations considered.

Table 5.4: Average prediction performance metrics under interpolation conditions.

Evaluation Window	**Metric**	**Weak Impurities**	**Product**	**Strong Impurities**
Process cycle at CSS	RMSE (mg/mL)	0.016	0.023	0.0033
	Accuracy (%)	93	98	95
Product collection window	RMSE (mg/mL)	0.000243	0.018	0.0028
	Accuracy (%)	99	98	99

5.3.3 Extrapolation

The extrapolation performance of the hybrid model is finally assessed by varying the feed concentration by ±20% relative to the nominal operating conditions. These scenarios extend beyond the training domain and test whether the model can generalise to larger variations in feed concentrations while maintaining predictive accuracy.

At reduced feed concentrations (Figure 5.7a), the amplitudes of all peaks predicted by the mechanistic model decrease. The hybrid model manages to follow that trend, reproducing the position and width of the species elution windows. More specifically, it is able to capture the elution profile of the product, showing however an underprediction during the product collection window, leading to lower predicted recovery. The weak impurities, by contrast, are slightly overpredicted, though their timing and return to baseline remain consistent with the mechanistic model. The strong impurities are underestimated in magnitude, yet their onset is still aligned with the mechanistic model, meaning that their small overlap with the product collection window is captured correctly. This results in a profile in which the relative sequencing of species is preserved, while the scaling of individual peaks deviates. This deviation produces a conservative bias, as the model underestimates product while not suppressing impurity contributions, avoiding any overestimation of predicted purity.

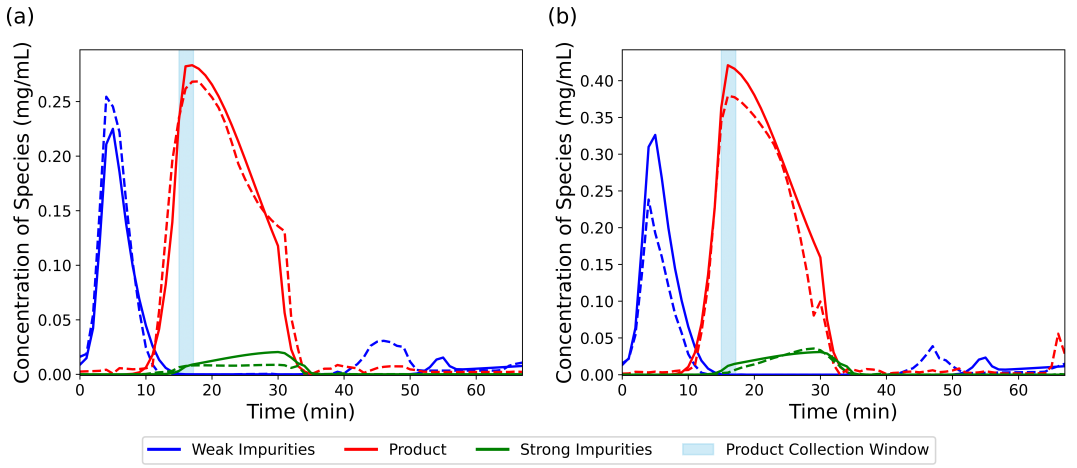


Figure 5.7: Comparison of outlet concentration profiles at CSS predicted by the mechanistic model (solid lines) and the hybrid model (dashed lines) under extrapolation conditions of feed variation: (a) −20% feed and (b) +20% feed.

At increased feed concentrations(Figure 5.7b), the respective amplitudes of all separation species rise. The hybrid again underpredicts the height of the product peak, while still managing to capture the timing of elution and shape of the peak. A similar profile can be observed for the strong impurities, which are underestimated in the onset of their elution, coinciding with the product collection window, but are aligned with the rest of the elution profile predicted by the mechanistic model. Finally, the weak impurities show the largest deviation from the mechanistic profile, underpredicting the height of the peak and shortening the elution window. As in the reduced-feed case, the hybrid model underestimates the product during the product collection window, leading to underestimation of the process KPIs, namely the product purity and the process yield.

This behaviour can be attributed to the structure of the hybrid model. By retaining the equilibrium isotherm, the model preserves the correct competitive dynamics between species, which explains why the timing of all elution windows remains consistent even outside the training range. The ANN, however, is responsible for scaling peak amplitudes, and it has only been trained on the narrower range of feed concentrations included in the training data. When asked to predict beyond this range, it underestimates the size of the product peak. This bias reflects the limited extrapolation capability of the network, which tends to compress responses at the boundaries of the training data (LeCun et al., 2012).

The performance metrics presented in Table 5.5 reflect this behaviour. Across the full cycle, the largest deviations arise from the underprediction of the elution peaks of the separation species. Within the product collection window, the errors are smaller in absolute terms but remain dominated by the underprediction of the product peak. Accuracy decreases compared to interpolation but remains above 90%, confirming that while deviations are more pronounced outside the training domain, the hybrid still captures the elution profiles of the species.

Table 5.5: Average prediction performance metrics under extrapolation conditions.

Evaluation Window	**Metric**	**Nominal Conditions**	**Interpolation**	**Extrapolation**
Process Cycle at CSS	RMSE (mg/mL)	0.029	0.033	0.0046
	Accuracy (%)	92	96	91
Product Collection Window	RMSE (mg/mL)	0.00023	0.035	0.0027
	Accuracy (%)	98	95	93

In practice, the direction of the error is significant. The hybrid systematically underestimates product within the collection window, which reduces yield predictions but does not risk overstating purity. Since purity is the hard constraint of the process, this type of conservative bias is less critical than the opposite case, where overestimation could lead to collection of product that does not meet the required purity specifications. The hybrid therefore is able to generalise under feed variations beyond the training domain, capturing the timing and shape of the elution peaks, while remaining conservative in its KPI predictions.

5.4 Conclusions

In this Chapter, a hybrid modelling framework was developed for the twin-column MCSGP process, combining a mechanistic description of competitive adsorption with an ANN for outlet concentrations at CSS. The resulting model reduced computational cost by 97% compared to the mechanistic model, while retaining sufficient accuracy for KPI evaluation.

Validation under nominal conditions confirmed that the hybrid model reproduces the mechanistic model with high accuracy, particularly within the product collection window where KPIs are evaluated. The interpolation tests showed that the model adapts reliably to variations in feed and flowrate within the training domain, capturing both amplitude scaling and shifts in elution timing while maintaining high fidelity. Under extrapolation, the hybrid preserved the timing and sequencing of elution windows but systematically underpredicted the product peak inside the product collection window, thus maintaining a conservative bias, where yield is underestimated, while purity is not overstated.

The findings demonstrate that hybrid modelling offers a computationally efficient and physically consistent surrogate for chromatographic separation. By retaining the essential dynamics of adsorption while accelerating predictions, the model provides a robust platform for repeated evaluation of process performance. The conservative bias observed under extrapolation further supports its use in decision-making contexts where purity must remain guaranteed. These features make the developed hybrid model particularly well suited for integration into online optimisation workflows, which will be explored in the following Chapter.

Chapter 6

Process optimisation

Outputs of this chapter have been disseminated in:

- **F. Michalopoulou** and M. M. Papathanasiou. Accelerated process optimization of chromatographic separations using a hybrid modeling approach. *IFAC-PapersOnLine*, 59(6):427–432, 2025.
- **F. Michalopoulou** and M. M. Papathanasiou. An approach to hybrid modelling in chromatographic separation processes. *Digital Chemical Engineering*, 14:100215, 2025. (Invited article in the *Emerging Stars* special issue)
- **F. Michalopoulou** and M. M. Papathanasiou. Assessment of data-driven modeling approaches for chromatographic separation processes. *AIChE Journal*, e18600, 2024. (Invited article in the *2024 Futures Issue*)

6.1 Introduction

Process optimisation lies at the core of chromatographic process development, where the competing objectives of product purity, yield, and productivity must be balanced under stringent regulatory and economic constraints (Kawajiri, 2021). Traditional approaches have relied on empirical tuning of gradients, switching points, and flowrates through extensive experimentation (Degerman et al., 2006; Natarajan and Lee, 2000). While such strategies can be effective at laboratory scale, they are inherently resource-intensive and limited in their ability to systematically explore the high-dimensional design space that governs separation performance (Chan et al., 2008; Close et al., 2014). The non-linear dynamics of adsorption, combined with tight operational constraints, make experimental trial-and-error poorly suited to the increasingly complex separation tasks encountered in biopharmaceutical manufacturing (Engell and Toumi, 2004).

The introduction of mechanistic modelling has established a rigorous framework for systematic optimisation in chromatography. First-principles descriptions of mass transfer and adsorption phenomena have been applied to design gradient programmes, determine optimal switching times, and quantify trade-offs between throughput, solvent consumption, and product recovery (Osberghaus et al., 2012; Teoh et al., 2001; Kawajiri and Biegler, 2008). These methods have also been extended to multicolumn configurations, where cyclic steady state (CSS) operation increases computational demands but enables substantial gains in yield and productivity (Taguado Menza et al., 2024). Collectively, these contributions have established mechanistic optimisation as a cornerstone of chromatographic process design, although the reliance on large-scale partial differential and algebraic equation (PDAE) systems can impose substantial computational costs that limit their application in iterative or real-time optimisation workflows.

As a result, surrogate modelling has been adopted to approximate high-fidelity models while retaining sufficient predictive power. Early applications integrated artificial neural networks (ANNs) into optimisation frameworks to accelerate the search for favourable elution profiles (Webb et al., 2009). More recent work has expanded the scope to multi-objective problems, simultaneously optimising cost of goods, yield, and purity across entire purification trains (Keulen et al., 2023; Liu and Papageorgiou, 2019). Gaussian process (GP) surrogates

have further enabled constraint-aware optimisation in ion-exchange chromatography, exploiting uncertainty quantification to guide adaptive sampling and improve convergence efficiency (Ding and Ierapetritou, 2024). Together, these approaches illustrate how surrogate-assisted optimisation extends the applicability of model-based design beyond the limits of direct mechanistic simulation.

More recently, hybrid models have been introduced into optimisation frameworks to combine mechanistic fidelity with the flexibility of neural networks. Early applications in preparative chromatography demonstrated how adsorption models could be replaced with hybrid surrogates to accelerate the optimisation of elution strategies (Nagrath et al., 2004). Later work extended this approach to capture and gradient operations, with equilibrium-constrained networks improving extrapolation under limited data and supporting the optimisation of operating conditions (Joshi et al., 2017; Narayanan et al., 2021b). Hybrid formulations have also been proposed for hydrophobic interaction chromatography, where they enabled parameter estimation and process optimisation (Ding et al., 2023). At larger scales, hybrid surrogates have been embedded into flowsheet optimisation frameworks to address multi-objective trade-offs in downstream design that include chromatographic steps (Pirrung et al., 2017). Collectively, these contributions highlight the potential of hybrid models to support chromatographic optimisation, though applications remain largely confined to aggregated performance indicators rather than the direct optimisation of full elution profiles.

In this Chapter, process optimisation of the twin-column multicolumn countercurrent solvent gradient purification (MCSGP) process is investigated using mechanistic, surrogate, and hybrid models. Building on the modelling frameworks introduced in Chapters 4 and 5, the developed models are embedded within optimisation workflows to evaluate convergence behaviour, computational requirements, and solution quality under strict purity constraints. Particular attention is given to the feasibility of surrogate-assisted optimisation under time-constrained settings, reflecting the requirements of online decision-making. The objective is to assess how model structure influences the efficiency and reliability of optimisation, and to demonstrate the extent to which surrogate-assisted approaches can provide a practical pathway towards real-time process optimisation.

6.2 Methodology

The optimisation framework developed in this Chapter is structured around three main components: (i) the formulation of the optimisation problem, (ii) the mathematical definition of the optimisation task, and (iii) the implementation of a Bayesian optimisation (BO) strategy. The methodologies underlying each of these components are discussed in the following sections.

6.2.1 Problem formulation

The optimisation tasks performed in this Chapter are carried out using the best-performing models developed in Chapters 4 and 5, namely the multiple-input/single-output (MISO) data-driven model employing the ReLU activation function and the isotherm-retaining hybrid model. These surrogates are employed as computationally efficient substitutes for the high-fidelity mechanistic process model, enabling repeated evaluations within an optimisation loop without the prohibitive burden of solving large systems of PDAEs. The central objective is to assess how such models can be embedded in an optimisation framework that maximises process yield while ensuring compliance with stringent purity requirements at CSS.

Both offline and online optimisation settings are considered. The offline case represents decisions made before process operation begins, when the feed composition has been measured and operating conditions can be set in advance. Two problem formulations are explored in this context. In the first, the optimisation is restricted to the three flowrates that govern the interconnected and batch phases of the twin-column MCSGP cycle. Flowrate exerts a direct influence on the residence time of solutes, the position and overlap of the elution peaks, and ultimately the balance between process yield and product purity. This restricted problem provides a baseline for assessing surrogate performance in a relatively low-dimensional setting, where the effect of flowrate manipulation on the yield–purity balance can be isolated. In the second scenario considered, the dimensionality of the problem is expanded to also include the initial modifier concentrations applied at the start of each phase. These concentrations shape the elution environment and strongly influence peak separation, making them candidates for offline optimisation alongside flowrates. Together, these two cases

assess the capability of the surrogates to solve optimisation problems of different dimensionality, while remaining consistent with operational parameters that are typically defined prior to execution.

In the online case, the scope is deliberately restricted to flowrate manipulation. Flowrates can be readily adjusted during operation, providing the most direct means of compensating for variability between feed batches. Other operational degrees of freedom, such as the switching times between the process steps, the imposed gradient profile, or the initial modifier concentration, remain fixed. This decision reflects both the scope of the trained surrogates, which were developed under fixed scheduling and gradient conditions, and the reality of process operation, where such parameters are typically defined during process development and are not readily adjustable in real time (Cabo-Calvet et al., 2014; Joshi et al., 2015). Feed composition and initial modifier concentration may vary between scenarios, but in the online case they are treated as external operating conditions rather than decision variables.

Across all scenarios, the objective is to maximise process yield while ensuring that product purity remains above the required purity specification throughout the fixed product collection window. Since the collection window is predetermined by the switching sequence between the various phases of the process, the optimisation is constrained to ensure the purity specification across the entirety of phase B1 (Figure 3.1). Purity is therefore enforced as a constraint, while yield serves as the measure of process efficiency.

This layered formulation, progressing from lower-dimensional flowrate optimisation to higher-dimensional cases incorporating initial modifier concentrations, ensures that the optimisation explores conditions that are both representative of the physical system and realistic for online implementation, while also testing the capability of the surrogate models to handle higher-dimensional problems in the offline setting. Applying the same problem definition across both the data-driven and hybrid surrogates, as well as the mechanistic model, allows differences in performance to be attributed directly to model structure, rather than to variations in scope. In this way, their efficiency in maximising process yield at fixed product purity, as well as their robustness across changes in feed and initial modifier conditions, can be systematically assessed.

6.2.2 Optimisation framework

The optimisation problem is expressed in the terms of the selected decision variables, namely the flowrates of the interconnected and batch phases, and, in the extended offline case, the initial modifier concentrations. These are treated as the control variables of the problem and are bounded within the validation range of the mechanistic model (Table 6.1). No other parameters of the cycle, such as switching times or the gradient profile, are subject to optimisation.

Table 6.1: Optimisation bounds for decision variables.

Input	Lower bound	Upper bound
Flowrates		
Flowrate during phases B1 and B2 (mL/min), (Q_B)	0.1	1
Flowrate during phases I1 (mL/min), (Q_{I1})	0.1	1
Flowrate during phases I2 (mL/min), (Q_{I2})	0.1	1
Modifier concentrations		
Feed concentration of modifier (mg/mL), ($C_{\text{feed}M}$)	1	3
Initial modifier concentration of Phases I1 for column 1 (mg/mL), (C_{M1})	2	3
Initial modifier concentration of Phases I1 for column 2 (mg/mL), (C_{M2})	0.8	1.2

The objective is to maximise process yield (Equation 3.15), while product purity (Equation 3.14) is enforced as a hard constraint across the fixed product collection window. Yield and purity are calculated from the predicted outlet concentrations during the product collection windows. In this case, productivity is not considered as an objective due to the fixed duration of the cycle, and yield is therefore adopted as a direct measure of separation performance. The optimisation problem is mathematically formulated as the minimisation of the negative yield, and can be expressed as:

$$\begin{aligned} \min_{u} \quad & F = -Y_j \\ \text{s.t.} \quad & u_{\min} \leq u \leq u_{\max} \\ & Pur_{av,j} \geq Pur_{\text{thresh}} \end{aligned} \tag{6.1}$$

where, Pur_{thresh} denotes the predefined purity specification, and u denotes the vector of decision variables, comprising the flowrates Q_B, Q_{I1}, Q_{I2} and, when included, the initial modifier concentrations $C_{\text{feed}M}$, C_{M1}, and C_{M2}. The bounds $u_{\min}$ and $u_{\max}$ for each decision variable are shown in Table 6.1. The mathematical formulation is identical across the offline and online settings, with only the dimensionality of u varying depending on the scope of the optimisation.

6.2.3 Bayesian optimisation

To solve this optimisation problem, a BO framework is employed. BO is particularly suited to problems where the objective function can only be assessed through costly evaluations, such as simulations or experiments, and where analytical gradients are not available to guide the search (Jones et al., 1998; Shahriari et al., 2016). In the present case, every candidate solution requires the simulation of a full twin-column MCSGP cycle at CSS to compute process yield and product purity, making the number of feasible evaluations limited. Under these circumstances, heuristic search methods or gradient-based solvers are unsuitable, whereas BO explicitly aims to minimise the number of evaluations by constructing a probabilistic surrogate of the objective function and updating it iteratively as new data become available (Figure 6.1).

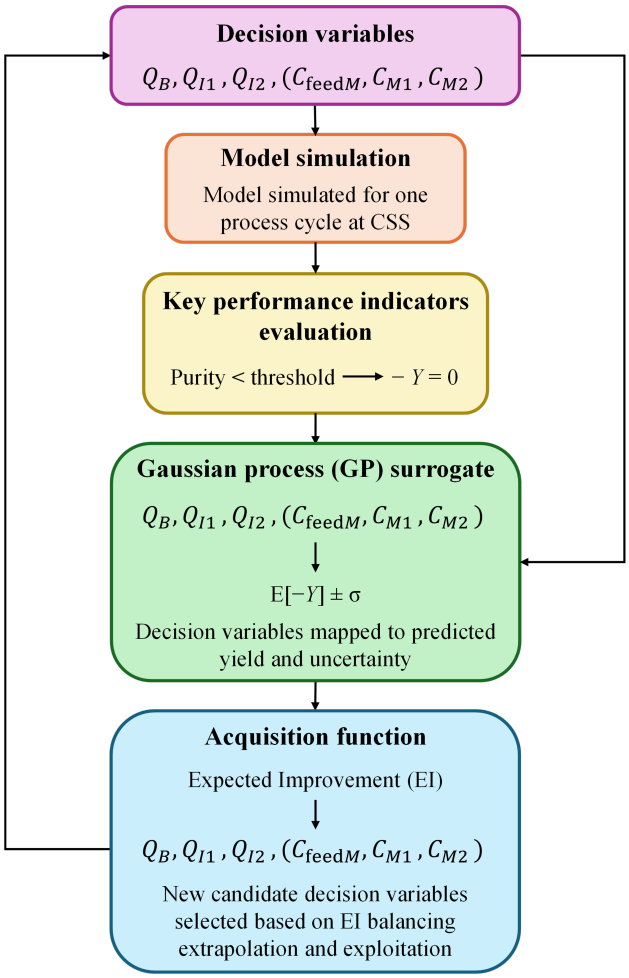


Figure 6.1: Optimisation workflow.

At the core of BO lies the GP prior, which models the objective function as a distribution over possible functions rather than a single deterministic mapping (Rasmussen and Williams, 2005). The GP provides both a prediction of the expected objective value at any candidate point in the design space and, crucially, an estimate of the uncertainty associated with that prediction (Snoek et al., 2012). This dual information is particularly important in the current setting, as the yield–purity trade-off is shaped by non-linear interactions between flowrates, modifier concentrations, and the separation dynamics. Sharp trade-off regions can exist where small changes in operating conditions result in pronounced shifts in feasibility. The GP framework captures these behaviours by adapting its variance locally, assigning high uncertainty in unexplored regions and reduced uncertainty where evaluations have already been performed.

The decision of where to evaluate next is determined by the acquisition function, which translates the GP's predictive mean and variance into a scalar utility. In this work, Expected Improvement (EI) is employed, as it has been shown to provide a reliable balance between exploration and exploitation (Jones et al., 1998). The exploitation aspect arises from favouring conditions expected to result in higher yield within the purity constraint, while exploration is maintained by assigning value to regions of high uncertainty (Brochu et al., 2010). This prevents the optimisation from prematurely focusing on a local optimum and ensures that the broader feasible domain is adequately mapped. In practice, EI identifies candidate operating conditions where either the predicted yield is already high or where the uncertainty is large enough that improvement over the incumbent best solution remains plausible.

The optimisation proceeds iteratively. For each set of decision variables selected by the acquisition function, the process models of the twin-column MCSGP system simulate one full cycle at CSS. Process yield and product purity are then computed over the fixed collection window, providing the key performance indicators (KPIs) used to evaluate the candidate. If purity falls below the predefined purity threshold, the candidate is flagged as infeasible, and the GP is updated accordingly to reflect this boundary (Gardner et al., 2014). By contrast, feasible solutions are incorporated into the GP posterior, refining the optimiser's understanding of both the objective surface and the feasible region. Over successive iterations, this process drives the search towards high-yield conditions that consistently satisfy the purity constraint, while avoiding unnecessary evaluations

in regions unlikely to produce viable results.

Finally, the optimisation objective is expressed as the minimisation of negative yield to align with the BO framework, which operates in minimisation form (Shahriari et al., 2016). While the task remains the maximisation of yield, this transformation ensures direct compatibility with the GP surrogate and the acquisition function. The optimisation thus proceeds iteratively, updating its internal surrogate of the objective function and refining its estimate of the feasible region until convergence is reached (Figure 6.1). The framework is identical across offline and online cases, with only the dimensionality of the decision vector varying depending on whether modifier concentrations are included alongside flowrates.

6.3 Results & discussion

6.3.1 Offline optimisation

The optimisation results are presented in two stages. The first stage considers offline optimisation, where the objective is to identify operating points that maximise yield while satisfying the purity constraint under fixed feed conditions. This setting provides a clear benchmark to compare the performance of the hybrid and data-driven models against the mechanistic model, since optimisation exposes not only the accuracy of the predicted profiles but also their ability to guide the search towards feasible and stable operating points. Results are shown first for the nominal feed composition and then for interpolation and extrapolation scenarios that test the ability of each surrogate model to operate reliably under domain shift. In all cases, optimisation is performed using the corresponding mechanistic, hybrid, or data-driven model, while the resulting operating points are evaluated by re-simulating them with the full mechanistic model to compute yield and verify purity.

6.3.1.1 Optimisation under nominal conditions

The BO framework is first applied to the optimisation of the phase flowrates under nominal feed conditions. The convergence of the three types of models, namely the mechanistic, the hybrid, and the data-driven, is shown in Figure 6.2. At nominal operation the yield reaches 85%, which provides a baseline for assessing improvements.

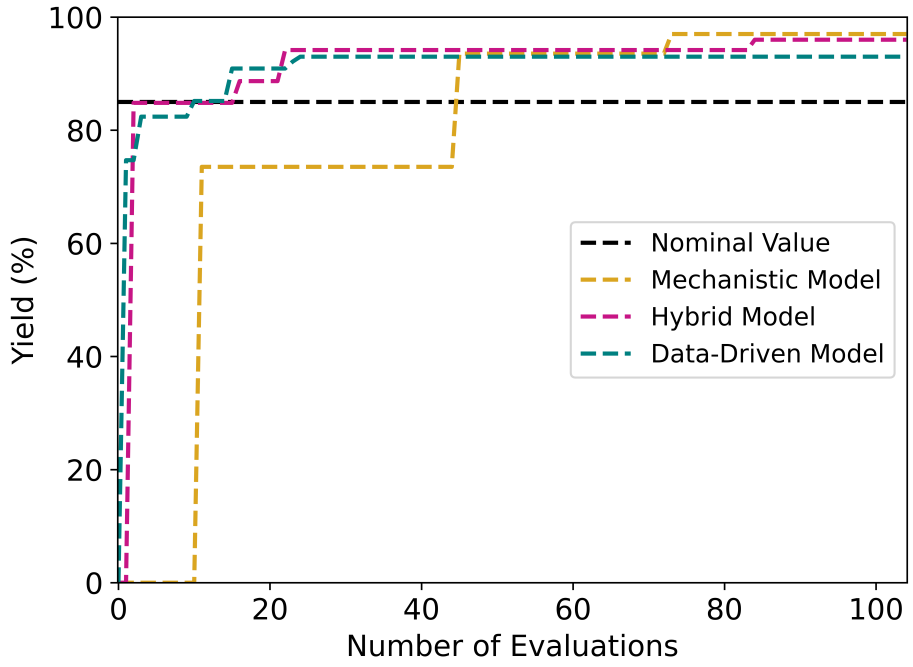


Figure 6.2: Convergence of Bayesian optimisation for offline optimisation of flowrates.

In this reduced problem, all models improve upon nominal performance while satisfying the purity specification, but differences arise in the final yields achieved and in the computational effort required. The mechanistic optimisation attains the highest yield at 97%, the hybrid model converges to a closely matching solution at 96%, and the data-driven model reaches 93% (Figure 6.2). These outcomes reflect the level of process knowledge embedded in each model. The full mechanistic description provides the most accurate representation of the system, the hybrid retains essential mechanistic structure alongside learned components, and the data-driven model relies solely on data without physics. As a result, the surrogates approximate the mechanistic optimum closely, with the hybrid tracking it more tightly than the purely data-driven model.

Another difference lies in the number of evaluations required to reach the respective optimal points, and the respective computational cost. The data-driven model reaches the optimum in just over 20 evaluations, while the hybrid model requires more than 80, and the mechanistic model more than 75. When these evaluation counts are translated into real process time these evaluation counts reveal the practical significance of the efficiency differences. The mechanistic optimisation requires several hours of computation, whereas the surrogates deliver their solutions within minutes, albeit with a small reduction in yield relative to

the mechanistic benchmark. Table 6.2 summarises these results, highlighting the trade-off between solution quality and computational cost.

Table 6.2: Offline Bayesian optimisation outputs.

	Mechanistic	**Hybrid**	**Data-driven**
Q_B (mL/min)	0.61	0.53	0.56
Q_{I1} (mL/min)	0.65	0.63	0.68
Q_{I2} (mL/min)	0.40	0.42	0.34
Purity (%)	98	98	98
Yield (%)	97	96	93
Computational time (h)	3.3	0.1	0.1

Small deviations are observed in the optimal flowrates selected by each model, which account for the different process yields achieved (Table 6.2). Such differences are expected when surrogates are used within BO, since they approximate rather than exactly reproduce the relationship between decision variables and yield. The optimiser may therefore settle at operating points that lie slightly offset from the mechanistic optimum. The probabilistic nature of BO, which balances exploration and exploitation, reinforces this effect. In practice these deviations translate into the small but systematic differences seen in the final yields, with the hybrid solution closest to the mechanistic benchmark and the data-driven solution slightly lower, while purity remains satisfied in all cases.

The dimensionality of the problem is then expanded to include the initial modifier concentrations as additional decision variables. This extended setting increases the complexity of the objective problem and accentuates the nonlinearities in the yield–purity trade-off. The mechanistic optimisation reaches 99% yield, the hybrid model converges to 97%, and the data-driven model to 93% (Figure 6.3). The efficiency differences are magnified in this higher-dimensional space. The hybrid and data-driven models converge in less than fifty evaluations, maintaining robust performance despite the increased dimensionality. In contrast, the mechanistic trajectory shows a slower progression. Its convergence shows only incremental yield gains during the first half of the campaign. Substantial improvement occurs only after two hundred iterations, with a gradual approach to the optimum, while a total of more than three hundred evaluations are needed before the mechanistic model reaches the maximum yield.

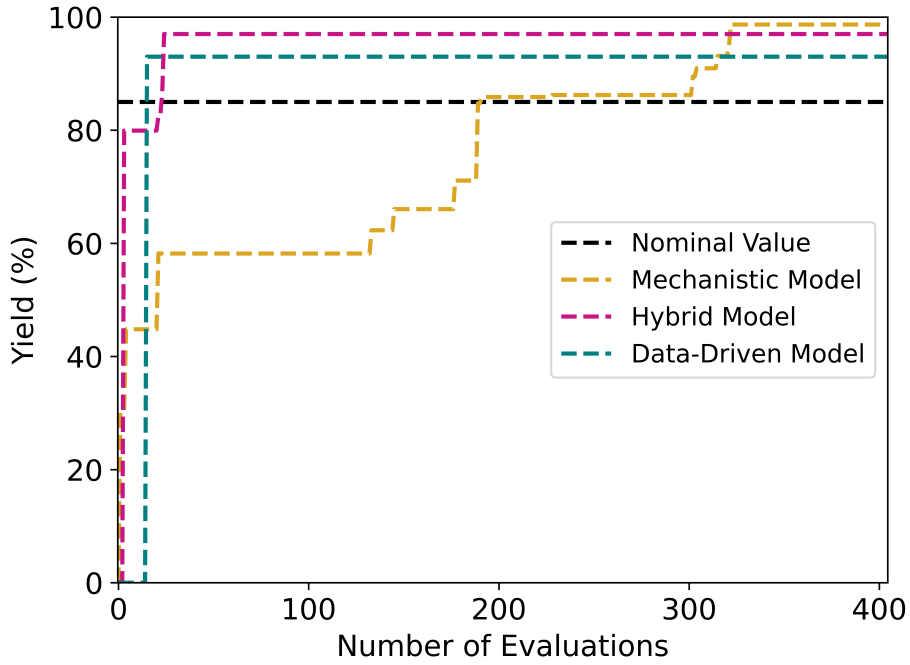


Figure 6.3: Convergence of Bayesian optimisation for offline optimisation of flowrates and initial modifier concentrations.

As in the flowrate optimisation case, the operating points identified by the surrogates exhibit small deviations from the mechanistic solution (Table 6.3). Examination of the optimal decision variables shows that the hybrid remained close to the mechanistic strategy, recovering 97% yield but with a higher Q_{I2} that slightly altered the recycling balance. The data-driven solution achieved only 93% yield and diverged more from the mechanistic structure, with lower Q_B and Q_{I1} and a reduced C_{M1}. These deviations explain the larger shortfall in yield and illustrate that the hybrid provides a closer compromise between fidelity and efficiency. The differences are consistent with surrogate approximation error and with the stochastic search inherent to BO, and they account for the slight shortfall in yield relative to the mechanistic benchmark. The difference lies in the resources required to obtain the solutions. The mechanistic optimisation requires several hundred simulations, each lasting two minutes, and therefore accumulates a computational time cost of more than ten hours. Such overhead renders the approach infeasible for any setting in which decisions must be made in real time. By comparison, the hybrid and data-driven models deliver improved, compared to nominal operation, operating points within minutes, making them suitable candidates for integration into online optimisation workflows.

Table 6.3: Extended offline Bayesian optimisation outputs.

	Mechanistic	**Hybrid**	**Data-driven**
Q_B (mL/min)	0.91	1.0	0.83
Q_{I1} (mL/min)	0.63	0.68	0.41
Q_{I2} (mL/min)	0.24	0.53	0.32
$C_{\text{feed}M}$ (mg/mL)	2.18	2.66	2.53
C_{M1} (mg/mL)	2.43	2.35	1.62
C_{M2} (mg/mL)	1.18	0.80	1.20
Purity (%)	98	98	98
Yield (%)	99	97	93
Computational time (h)	11.7	0.4	0.2

6.3.1.2 Optimisation under variable feed conditions

To assess the robustness of the hybrid and data-driven surrogates beyond the nominal case, BO is repeated for two additional feed scenarios (Table 6.4). The first corresponds to an interpolation point, located within the training domains of both models but distinct from the operating points used during training, where the feed concentrations of the weak impurities ($C_{\text{feed}W}$), product ($C_{\text{feed}P}$), and strong impurities ($C_{\text{feed}S}$) are decreased by 10% from the nominal conditions. The second corresponds to an extrapolation point, chosen deliberately within the training range of the data-driven model but outside that of the hybrid, where feed concentrations are increased by 20% compared to nominal operation. These scenarios test whether the predictive advantages of the hybrid observed in Chapter 5 extend to optimisation and illustrate how each surrogate behaves when driving the search for feasible operating points under domain shift.

Table 6.4: Feed concentrations for Bayesian optimisation under interpolation and extrapolation conditions.

Scenario	$C_{\mathbf{feed}W}$ **(mg/mL)**	$C_{\mathbf{feed}P}$ **(mg/mL)**	$C_{\mathbf{feed}S}$ **(mg/mL)**
Interpolation	0.065	0.35	0.035
Extrapolation	0.08	0.5	0.05

In the interpolation case, all three models identified feasible operating conditions that met the 98% purity constraint, while the mechanistic model converged

to a maximum yield of 97% (Table 6.5). The hybrid model reproduced the mechanistic outcome both in terms of yield and flowrate structure, with Q_B and Q_{I2} essentially unchanged and Q_{I1} only marginally lower. The data-driven model achieved a slightly reduced yield of 95% and reached this point through a redistribution of the internal flows, lowering Q_{I1} to 0.25 mL/min and increasing Q_{I2} to 0.71 mL/min. This adjustment altered the balance of material recycling and led to small but systematic product losses at the collection window. Although both surrogates converged to feasible solutions, the hybrid remained closer to the mechanistic operating strategy, indicating that the embedded adsorption equations stabilised the optimisation trajectory and limited deviations that reduced product recovery. The interpolation case therefore reflected the behaviour already observed for the nominal optimisation, with both the hybrid and the data-driven models performing well.

Table 6.5: Bayesian optimisation outputs under interpolation and extrapolation conditions.

Scenario	**Model**	Q_B **(mL/min)**	Q_{I1} **(mL/min)**	Q_{I2} **(mL/min)**	**Purity (%)**	**Yield (%)**
Interpolation	Mechanistic	0.60	0.36	0.65	98	97
	Hybrid	0.66	0.33	0.65	98	97
	Data-driven	0.65	0.25	0.71	98	95
Extrapolation	Mechanistic	0.70	0.58	0.41	98	95
	Hybrid	0.31	0.43	0.56	98	86
	Data-driven	0.47.	0.53	0.46	98	92

The extrapolation case exposed clearer differences between the models (Table 6.5). The mechanistic optimum delivered a yield of 95% at 98% purity with a solution biased towards higher Q_{I1}, reflecting the need for greater recycling of weak impurity and product overlapping fractions at elevated feed concentrations. The data-driven model, trained explicitly in this region, reproduced this trend, selecting $Q_{I1} = 0.53$ mL/min and $Q_{I2} = 0.46$ mL/min to achieve 92% yield. The hybrid, operating outside its training domain, departed more substantially from the mechanistic structure. Its solution reduced Q_B to 0.31 mL/min while balancing Q_{I1} and Q_{I2} at intermediate values of 0.43 and 0.56 mL/min. This choice preserved purity but limited throughput and altered the recycling pattern in a way that reduced the recoverable product compared to the other models, resulting in a yield of 86%. The outcome illustrates the limits of extrapolation. The embedded isotherm ensured physically consistent behaviour, yet the lack of training data in this region meant that the surrogate did not capture the

correct scaling of yield with increased feed load. Rather than converging to an infeasible point, the hybrid converged to a relatively more conservative solution that guaranteed product purity at the expense of process yield.

These results show that the hybrid model combines predictive accuracy within its training domain with resilience under extrapolation. In the interpolation case the hybrid reproduced the mechanistic optimum both in yield and operating strategy, whereas the data-driven model achieved slightly lower process yield. In the extrapolation case the hybrid maintained feasibility and physical consistency but converged to a lower yield compared to the data-driven model, reflecting the absence of training coverage in this region. This behaviour demonstrates that the hybrid extrapolates conservatively in optimisation, prioritising product purity. This robustness is particularly relevant for process development, where feed variability is common and models must remain reliable under domain shifts, while purely data-driven surrogates depend more heavily on broad training coverage to maintain feasibility.

6.3.2 Online optimisation

The second stage of the results focuses on time-constrained optimisation, where the surrogate models are embedded in online strategies that update operating decisions during process execution. Building on the offline results, these scenarios are more representative of practical operation, where solutions must be obtained under strict time limitations and in the presence of variability across cycles. Two levels of time resolution are examined, with cycle-to-cycle optimisation updating conditions once per process cycle and minute-by-minute optimisation adapting conditions continuously within a cycle. As in the offline case, the corresponding mechanistic, hybrid, and data-driven models are used to propose operating decisions within the optimisation loop, while final performance metrics and elution profiles are obtained by re-simulating the optimised conditions using the mechanistic model.

6.3.2.1 Cycle-to-cycle optimisation

To evaluate the performance of the developed models online, a cycle-to-cycle optimisation scenario is considered under nominal feed conditions, where the search for optimal flowrates is constrained to one hour, corresponding to the approxi-

mate duration of a single twin-column MCSGP cycle. This time limit reflects a realistic window for implementing optimisation between consecutive runs. The objective remains to maximise yield under the fixed purity specification, with the optimiser allowed to explore only as many iterations as can be completed within the imposed horizon.

Within the one-hour horizon, both surrogate models converge to the same respective operating points as in the offline optimisation, reaching a yield of 96% with the hybrid model and 93% with the data-driven, while maintaining the purity specification throughout (Table 6.6). The data-driven model attains this optimum in less than one minute, while the hybrid requires longer but still converges well within the available time. The trajectories are consistent across runs. Yield increases steadily as the optimiser adjusts the flowrates, and any improvement in yield is considered only when the purity constraint is met. This is particularly important for cycle-to-cycle use, since it demonstrates that surrogate-assisted optimisation can provide a reliable update to the operating conditions at the end of each cycle, ensuring that the product specifications will be met, while still improving process performance compared to nominal operation.

Table 6.6: Cycle-to-cycle Bayesian optimisation outputs.

	Mechanistic	**Hybrid**	**Data-driven**
Q_B (mL/min)	0.52	0.53	0.56
Q_{I1} (mL/min)	0.76	0.63	0.68
Q_{I2} (mL/min)	0.22	0.42	0.34
Purity (%)	98	98	98
Yield (%)	79	96	93
Evaluations	30	1200	6000

The mechanistic model, by contrast, is fundamentally limited by its computational cost. With each evaluation requiring approximately two minutes, fewer than 30 iterations can be executed within the one-hour horizon. This restricts the optimiser to only a small fraction of the search space, leaving it unable to fully explore yield–purity trade-offs. The result is that the mechanistic optimisation plateaus below 80% yield and does not reach the optimum (Table 6.6). In practice, the optimiser cannot escape suboptimal regions because the evaluation budget is exhausted before meaningful improvements are achieved. This outcome

highlights why, even though the mechanistic model remains the most accurate representation of the process, its direct use is infeasible when optimisation must be carried out between consecutive cycles.

The difference between model performance is also visible in the outlet concentration profiles generated at the optimised operating points (Figure 6.4). The hybrid surrogate identifies flowrates that recover a greater fraction of the product peak within the collection window while maintaining purity above the threshold, resulting in visibly higher product recovery compared to the data-driven model. The data-driven surrogate also improves recovery, but some product remains outside the window, consistent with its slightly lower reported yield.

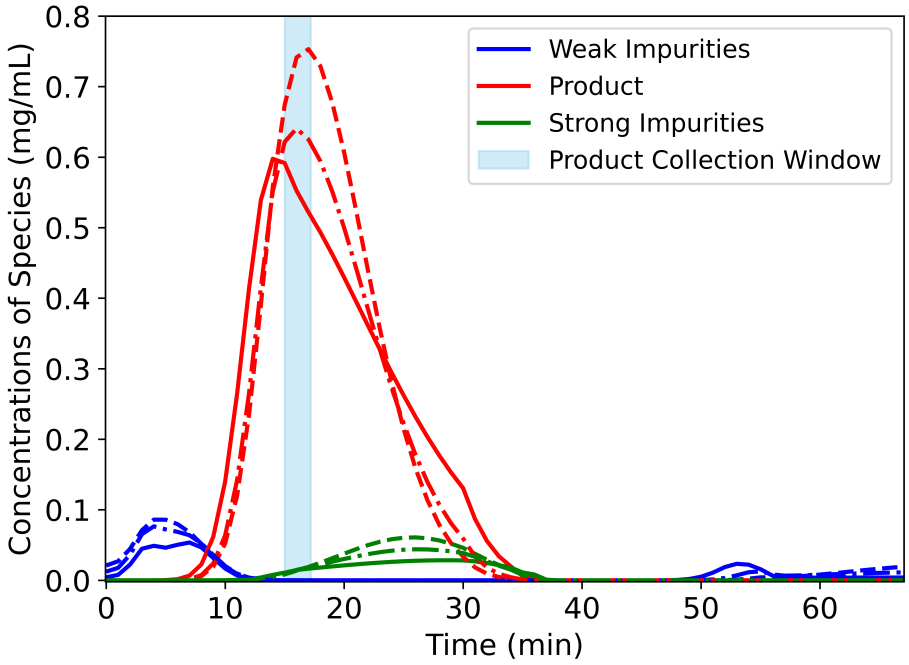


Figure 6.4: Comparison of the elution profiles of the weak impurities (blue), product (red), and strong impurities (green) at CSS, as optimised by the mechanistic (solid lines), hybrid (dashed lines) and data-driven (dash-dot lines) models in the cycle-to-cycle optimisation scenario.

By contrast, the mechanistic optimisation fails to refine the flowrates adequately within the one-hour horizon. As a result, the product peak extends beyond the collection window, causing part of the product to be recycled rather than collected, even though the purity specification is met. This highlights that the mechanistic model is not only computationally impractical under time-constrained settings but also operationally suboptimal, as the limited search prevents it from exploiting the full recovery potential of the process. In this

context, the hybrid model provides the best compromise between computational feasibility and solution quality, while the data-driven models prioritises speed and still achieves a high yield, albeit slightly below the hybrid.

6.3.2.2 Minute-by-minute optimisation

To test the feasibility of even finer time-scale optimisation, a minute-by-minute scenario is considered under nominal feed conditions. In this setting, the optimiser is given a time horizon of one minute to identify improved operating conditions. This represents the most stringent case, where optimisation must be performed in real time alongside process operation.

Under these conditions, the mechanistic model is rendered unusable (Table 6.7). Since a single evaluation requires approximately two minutes, no candidate solution can be completed within the imposed horizon. The model is therefore unable to contribute even a single evaluation, which makes optimisation impossible. This confirms that while the mechanistic description provides the most accurate representation of the system, its computational cost excludes it entirely from use in a minute-by-minute online optimisation context.

Table 6.7: Minute-by-minute Bayesian optimisation outputs.

	Mechanistic	**Hybrid**	**Data-driven**
Q_B (mL/min)	-	0.50	0.56
Q_{I1} (mL/min)	-	0.72	0.68
Q_{I2} (mL/min)	-	0.30	0.34
Purity (%)	-	98	98
Yield (%)	-	88	93
Evaluations	-	20	100

The hybrid model, although significantly faster, is still constrained by the imposed horizon. Each evaluation requires approximately three seconds, which allows only a limited number of iterations to be completed before the horizon is reached. The optimiser is able to explore feasible operating points and deliver improvements over nominal operation, but it does not converge fully to the offline optimum. Yields remain below the 96% achieved in the cycle-to-cycle and offline cases, highlighting how the available time budget constrains optimisation performance. By contrast, the data-driven model is able to complete one hundred

iterations within the one-minute horizon, owing to its structural simplicity. This enables the optimiser to explore the design space extensively, refine its search, and once again converge to the same optimum identified in the offline and cycle-to-cycle scenarios. As such, process yield of 93% is achieved while product purity remains satisfied throughout the product collection window.

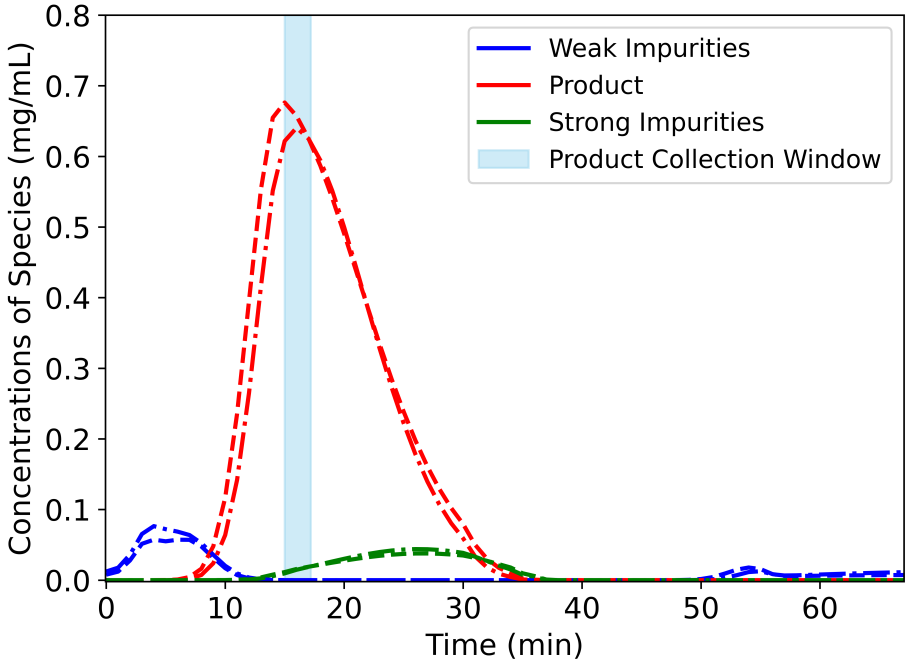


Figure 6.5: Comparison of the elution profiles of the weak impurities (blue), product (red), and strong impurities (green) at CSS, as optimised by the hybrid (dashed lines) and data-driven (dash-dot lines) models in the minute-by-minute optimisation scenario.

The outlet concentration profiles illustrate these differences directly (Figure 6.5). The data-driven model produces profiles that are identical to the offline optimum it had previously identified, with the collection window capturing nearly the entire product peak at the required purity. The hybrid model also improves product recovery, but the product peak remains less confined within the product collection window, leading to a reduced process yield of 88% compared to the cycle-to-cycle optimisation scenario. Finally, the absence of mechanistic profiles reflects the fact that the model cannot deliver even a single evaluation within the allotted one-minute horizon. These comparisons emphasise that only the data-driven surrogate is able to reproduce the optimal separation behaviour under real-time constraints, while the hybrid surrogate provides partial improvement but is limited by the reduced number of iterations it can complete.

6.4 Conclusions

In this Chapter, BO was applied to the twin-column MCSGP process using mechanistic, hybrid, and data-driven models under strict purity constraints. The framework was evaluated across offline, cycle-to-cycle, and minute-by-minute optimisation scenarios, allowing the influence of model structure and computational efficiency on optimisation performance to be systematically assessed.

The offline results established the mechanistic model as the benchmark, consistently delivering the highest yields but at the cost of long runtimes. The hybrid model reproduced these results with only marginal deviations when operating within its training bounds, while reducing computational cost by more than an order of magnitude, and maintained feasibility even when extrapolating to unknown feed compositions. The data-driven model provided further speed gains with a small reduction in yield. Under the one-hour cycle-to-cycle horizon, the mechanistic model was unable to converge to the optimum, whereas both surrogates maintained reliable performance. The hybrid model provided the closest match to the offline mechanistic benchmark, while the data-driven model delivered slightly lower yields but converged more rapidly. In the most stringent minute-by-minute case, only the data-driven model was capable of recovering the offline optimum, as the mechanistic model could not complete a single evaluation and the hybrid was limited by the restricted number of iterations.

Placed in the wider context of this thesis, these results highlight that optimisation in chromatography does not depend on a single universally best model, but on selecting the structure most suited to the application. The mechanistic model remains the reference for accuracy and is best used for process design, the hybrid model offers a balanced compromise for cycle-to-cycle optimisation, and the data-driven model enables optimisation under real-time constraints. Together, these findings underline the importance of a portfolio approach, where mechanistic, hybrid, and data-driven models are deployed in complementary roles to enable practical and efficient optimisation across different timescales of chromatographic operation. Building on this foundation, the next Chapter introduces a reinforcement learning (RL) framework for real-time control of the twin-column MCSGP process, extending beyond optimisation of fixed process schedules to adaptive, phase-aware operation under dynamic process conditions.

Chapter 7

Reinforcement learning for adaptive control

Outputs of this chapter have been disseminated in:

- **F. Michalopoulou**, Emmanouil Papadakis, and M. M. Papathanasiou. A reinforcement learning control strategy for chromatographic separation processes. *Submitted to Chemical Engineering Research and Design*, Under review. (Invited article in the *Emerging Stars II* special issue)
- **F. Michalopoulou** and M. M. Papathanasiou. Integrated dual adaptive control of continuous chromatographic separation processes via reinforcement learning. *Journal of Chromatography A*, 1768, 466638, 2026. (Invited article in the *Advances in Preparative Chromatography* special issue)

7.1 Introduction

The increasing complexity of biomanufacturing systems, alongside the necessity for improved efficiency and tighter product specifications, has led to a growing interest in digitalisation strategies that support real-time decision making (Pantelides and Pereira, 2024). Among these, advanced modelling and control tools have become key for dynamic optimisation, especially in processes with complex interactions, high variability, and strict performance requirements (Papathanasiou et al., 2019a). While mechanistic models have been playing a central role in process design and offline optimisation, their integration into real-time workflows remains challenging, in part due to the computational burden of solving detailed, high-dimensional systems in an online setting (Alhajeri et al., 2024; Grossmann et al., 2010).

As chromatographic separation continues to play a central role in downstream processing, the shift towards continuous and semi-continuous operation has led to increased interest in control strategies that can accommodate cyclic operation, frequent switching events, and time-varying input conditions (Xie et al., 2023). In such systems, early decisions within a cycle play a decisive role in shaping product quality and recovery, as their effects propagate through subsequent phases, thereby increasing the importance and complexity of the controller's task. While Process Analytical Technologies (PATs) offer avenues for online monitoring, their integration remains uneven, and data sparsity or delays often limit the design of feedback-driven control systems (Kim et al., 2021; Shastri et al., 2023).

Mechanistic models, typically derived from first principles and formulated as systems of partial differential and algebraic equations (PDAEs), have been widely used to describe the mass transfer and adsorption phenomena underlying chromatographic behaviour (Shekhawat and Rathore, 2019). These models support design and optimisation tasks and have formed the basis for advanced control strategies, such as model predictive control (MPC) (Suvarov et al., 2012). However, their computational demands, particularly when simulating spatially distributed dynamics, can often act as a barrier to their deployment in closed-loop control, especially under uncertainty or in the presence of fast disturbances (Eslami and Jungbauer, 2024). In such settings, implementation often requires reduced-order models, fast solvers, and reliable state estimation, all of which can be challenging when working with high-dimensional distributed models and

partial observability (Papathanasiou et al., 2016). As a result, real-time MPC for chromatography can involve substantial simplification or rely on application-specific tuning.

Reinforcement learning (RL) offers a complementary approach to controller design, in which a policy is learned through repeated interaction with the process. Unlike conventional model-based methods, RL does not require online optimisation or an internal model at deployment (Buşoniu et al., 2018). Instead, by interacting with the environment and adjusting its policy based on cumulative reward, an RL agent can learn to make informed decisions under partial observability and in the presence of delayed effects (Devarakonda et al., 2025). While RL has been applied to a range of chemical engineering problems, including scheduling, energy systems, and batch process optimisation, its use in chromatographic control remains limited (Alhazmi et al., 2022; Petsagkourakis et al., 2020; Yifei and Lakshminarayanan, 2023). A small number of studies have explored its application in liquid chromatography, where column variability or non-linear dynamics present particular challenges to conventional strategies (Kensert et al., 2024; Niezen et al., 2025; Nikita et al., 2021). However, to date, there has been little work examining how RL can be used to control continuous or semi-continuous multicolumn systems from start-up, under realistic conditions involving uncertainty, switching events, and noisy measurements (Andersson et al., 2023).

In this Chapter, an RL controller for real-time control of continuous chromatographic separation is developed, focusing on the twin-column multicolumn countercurrent solvent gradient purification (MCSGP) process. The controller is trained using a high-fidelity mechanistic process model, which simulates the underlying physicochemical behaviour of the system and serves as the RL agent's environment. The RL agent observes process outputs and learns to adjust key control inputs, such as flowrates and switching times, based on observed system states. A reward function is defined to reflect process objectives including purity, yield, and operational efficiency. The trained agent is evaluated under a range of dynamic conditions, including variable feed composition and measurement noise. The results demonstrate that the RL controller can maintain process performance under variability and uncertainty, offering a flexible and adaptive alternative for chromatographic process control.

7.2 Methodology

The RL controller developed in this Chapter is designed following five main steps: (i) problem formulation, (ii) selection of control framework, (iii) identification of input-output structure, (iv) development of reward function, and (v) controller training. The methodologies underlying each of these steps are discussed in the following sections.

7.2.1 Problem formulation

The objective of this Chapter is to design a controller for the twin-column MCSGP process that operates in real time from process start-up, adjusting flowrates dynamically across different phases of the chromatographic cycle. The process structure and operating phases are described in detail in Chapter 3. Here, the focus is on formulating the control task from the perspective of process operation and measurement availability.

The twin-column MCSGP process is governed by cyclic transitions between interconnected and batch phases. The timing and operating conditions during these transitions affect component elution, product recovery, and impurity retention. As such, during operation, the controller must determine how fast to operate during each phase and, implicitly, when to switch to the next. In this formulation, flowrate serves as the primary manipulated variable, influencing both the elution profile and the timing of process steps. The controller is therefore expected to determine both the value of the flowrate and the duration for which it is applied before transitioning to the next phase. These control actions are not scheduled in advance but must be selected online based on process observations.

In contrast to the online optimisation framework adopted in Chapter 6, the control formulation considered here allows both flowrate regulation and the effective control of phase durations to be treated within a single decision-making strategy. Whereas online optimisation typically requires restricting the decision space to a limited set of independent variables to maintain real-time tractability, the present framework shifts the computational burden associated with higher-dimensional decision spaces to controller development rather than online execution. This enables flowrates and switching times to be treated as interdependent decisions within the controller's decision process, providing a more flexible rep-

resentation of the operating degrees of freedom.

To support timely yet effective decision-making, control actions are sampled at fixed one-minute intervals. This balances responsiveness with the physical dynamics of the system, ensuring that each adjustment has time to influence the process before the next decision is made. It also provides a practical basis for real-time deployment, avoiding the inefficiencies of overly frequent control updates while maintaining sufficient resolution for dynamic operation.

Online measurements are assumed to be obtained through standard PATs, such as UV detectors (Shastri et al., 2023). These provide aggregate elution signals, which correspond to the total outlet concentration, typically without resolving individual species. As such, the controller does not have access to internal column profiles or component-specific concentrations, nor to process key performance indicators (KPIs) such as purity or yield, which require offline measurements. Instead, the design relies on observable outputs that can be acquired and interpreted during operation, including elution patterns, switching triggers, and flowrate histories. Nevertheless, ensuring product purity remains the central objective of the control strategy. The controller must therefore learn to infer progress toward this objective from observable trends in outlet signals and use these to guide real-time decisions.

To be practically useful, the controller is intended for deployment from process start-up, without requiring pre-optimised switching schedules or a stabilisation period before on-spec operation is achieved. It is expected to interpret early signals and make informed decisions from the initial cycle onwards. In addition, to support flexible product quality specifications, the target purity threshold is treated as a tuneable supervisory input to the controller. This value can be set at the beginning of the process to reflect current production targets or regulatory requirements. The controller must then interpret this threshold during operation and adjust its decisions accordingly. This enables a single controller structure to accommodate varying product specifications without the need for redesign or manual retuning, supporting adaptive use across different scenarios.

Although the control logic is based primarily on observable outputs, the design concept incorporates the high-fidelity mechanistic model of the system as a potential auxiliary tool. This model is restructured to simulate 1 minute intervals at minimal computational cost, making it suitable for online use. As such, it can support the interpretation of noisy or ambiguous signals, estimate unmeasured

concentrations, or complement sensor outputs when direct measurements are unavailable. In this way, the model acts as a corrective mechanism that enhances observability without becoming a dependency for controller deployment.

Taken together, the control strategy is intended to operate in an adaptive and phase-aware manner, responding to time-varying conditions while maintaining responsiveness to the current process step. The objective is to develop a control policy that leverages observable elution patterns to drive real-time flowrate adjustments, indirectly governing both separation performance and transition timing across the full chromatographic cycle. To achieve this, an RL framework is adopted for controller development and training.

7.2.2 Control framework

The control task defined above is framed as a sequential decision-making problem, in which a control policy must be learned to map process observations to flowrate decisions at each step of the chromatographic cycle. This is formulated as a partially observable Markov decision process (POMDP), reflecting the fact that the controller cannot observe the full internal state of the system (Zhao et al., 2023). Instead, decisions must be based on partial, high-level output measurements obtained from the process during operation.

To solve this control problem, an RL framework is adopted, in which the controller (or agent) interacts with a simulated environment to learn a policy that maximises long-term process performance. At each decision point, the agent receives a set of observable outputs (S_t), selects a flowrate to apply (A_t), and receives a scalar reward (R_{t+1}) that encodes process objectives such as product purity, process yield, and operational efficiency (Figure 7.1). Over repeated episodes, the agent improves its policy by adjusting its decision-making strategy to maximise cumulative reward.

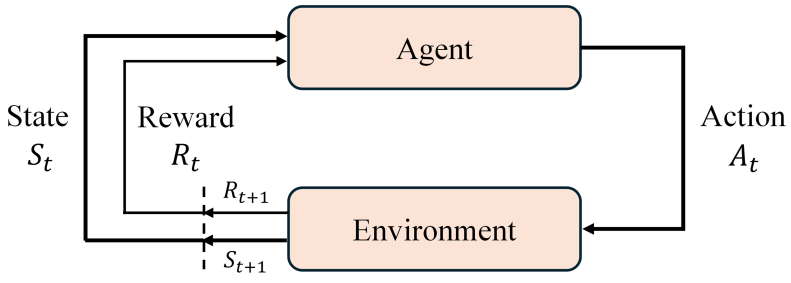


Figure 7.1: Reinforcement learning framework.

The controller is implemented using the Proximal Policy Optimisation (PPO) algorithm, a widely used policy-gradient method that combines sample efficiency with robust convergence (Schulman et al., 2017). PPO constrains the magnitude of policy updates to prevent large, destabilising changes during training, which makes it well suited for problems with continuous action spaces, partial observability, and sparse or noisy reward signals, which are frequently encountered in real-time chromatographic operation. This approach enables the agent to learn how to operate across all process phases, adapting its control decisions based on the current step, observed elution behaviour, and historical outcomes.

While the PPO algorithm is model-free, meaning that it does not require an internal model of the system to make decisions, the training process in this study is model-informed. The agent interacts with a simulated process environment governed by a high-fidelity mechanistic model of the twin-column MCSGP system. The environment provides state transitions and reward feedback based on the underlying physics of the system, thus allowing the agent to learn a policy that is both data-driven and physically grounded.

7.2.3 Input-output structure

The RL agent interacts with the chromatographic system through a structured set of inputs and outputs that reflect the real-time constraints of the process. The observation space comprises six features, each selected to convey critical information about the current process state as inferred from available measurements. These include both directly observable outputs and process-level parameters that shape the controller's response (Table 7.1).

Table 7.1: Controller inputs.

Input	**Description**	**Source**	**Range**
Step index	Encodes current process phase	Internal logic	[0.0, 0.875]
Weak impurities concentration	First impurities at outlet	PAT + auxiliary model	[0, 1] mg/mL
Product concentration	Target component at outlet	PAT + auxiliary model	[0, 1] mg/mL
Strong impurities concentration	Second impurities at outlet	PAT + auxiliary model	[0, 1] mg/mL
Flowrate	Previous control input	Environment history	[0.1, 1.0] mL/min
Purity threshold	Target specification	Supervisory input	[0.95, 0.98]

The first element of the observation vector encodes the current process step, represented as a normalised index ranging from 0 to 0.875 across the eight-phase cycle. This provides phase-awareness to the agent, enabling it to interpret

signals in the context of the ongoing separation task. The next three inputs comprise the eluted concentrations of the weak impurities, the target product, and the strong impurities at the elution column outlet. These are estimated from aggregate signals provided by online PAT instruments (such as UV detectors) and, when necessary, de-noised or complemented using the auxiliary mechanistic model. The fifth input is the previously applied flowrate, included to capture the influence of past control actions on current elution behaviour and to help the agent learn process dynamics over time. Finally, the purity threshold is included as a tunable supervisory input that remains fixed throughout each episode. This threshold reflects the desired product quality for a given process scenario and allows the same policy to be deployed flexibly across different operating targets without retraining.

All the inputs are rescaled prior to being passed to the policy network, ensuring numerical stability and improving generalisation during training. This normalisation also supports consistent performance across a range of operating scenarios, including those involving variations in feed concentrations and column loading.

The controller's output is a single continuous action representing the flowrate to be applied to the eluting (right-hand) column at the current timestep. This flowrate is bounded between 0.1 and 1.0 mL/min, in accordance with physical equipment limitations. It should be noted, that the agent is not responsible for adjusting the flowrate of the left-hand column, which is executing the recycling and feeding tasks (Figure 3.1). This column is operated at its maximum permissible flowrate (1.0 mL/min) to guarantee effective recycling and separation of the components.

7.2.4 Reward function

The reward function serves as the principal mechanism through which the RL controller evaluates its performance and adapts its behaviour to maximise long-term process objectives. In the context of twin-column MCSGP operation, these objectives include the production of high-purity product fractions, minimisation of impurity recycling, and efficient use of operating time. To this end, the reward function is constructed to reflect both the separation goals and the structural constraints of the process environment, enabling the controller to adjust the

elution column flowrate in a manner that improves overall performance over time. Given the cyclic nature of the process, the reward function is modular and phase-dependent, with each chromatographic phase associated with distinct operational priorities and control behaviours.

The twin-column MCSGP cycle is composed of eight discrete steps, which are symmetrically divided into two half-cycles (Figure 3.1). However, from a control perspective, each half-cycle is governed by an identical sequence of operational phases. Thus, the reward formulation is presented in terms of four functional phases, which are repeated in both halves of the cycle.

Each of the system phases presents a distinct trade-off between product retention, impurity removal, and process efficiency. The reward function is designed to reflect these trade-offs through weighted contributions of component concentrations and time-dependent penalties. The general form of the reward function at timestep t is given by:

$$r_t = Q_t \cdot f(C_{Wt}, C_{Pt}, C_{St}) - \delta \cdot t \tag{7.1}$$

where the structure of the function f varies by phase, Q_t is the flowrate at time t, C_{Wt}, C_{Pt}, and C_{St} the outlet concentrations of the weak impurities, the product, and the strong impurities respectively at time t, and δ the step duration penalising constant. The specific reward expressions for each phase are:

- Phase I1 - Recycling of product with weak impurities:

$$r_{t,\mathrm{I1}} = Q_t \cdot (\alpha_{\mathrm{I1}} \cdot C_{Wt} - \beta_{\mathrm{I1}} \cdot C_{Pt}) - \delta_{\mathrm{I1}} \cdot t \tag{7.2}$$

- Phase B1 - Product collection window:

$$r_{t,\mathrm{B1}} = \beta_{\mathrm{B1}} \cdot Q_t \cdot C_{Pt} - \delta_{\mathrm{B1}} \cdot t \tag{7.3}$$

- Phase I2 - Recycling of product with strong impurities:

$$r_{t,\mathrm{I2}} = Q_t \cdot (-\beta_{\mathrm{I2}} \cdot C_{Pt} + \gamma_{\mathrm{I2}} \cdot C_{St}) - \delta_{\mathrm{I2}} \cdot t \tag{7.4}$$

- Phase B2 - Impurities elution window:

$$r_{t,\mathrm{B2}} = Q_t \cdot (\alpha_{\mathrm{B2}} \cdot C_{Wt} - \beta_{\mathrm{B2}} \cdot C_{Pt} + \gamma_{\mathrm{B2}} \cdot C_{St}) - \delta_{\mathrm{B2}} \cdot t \tag{7.5}$$

The constants α, β, and γ are positive weights that encode the relative importance of component-specific concentrations in each phase, namely of the weak impurities, the product, and the strong impurities respectively. Their values are chosen based on process knowledge and are summarised in Table 7.2. For example, low values of β in the recycling phases (I1 and I2) penalise high product concentration and therefore insufficient separation, while still rewarding the recycling of the stream through higher values of α and γ. Similarly, during the elution phase (B1), β incentivises maximum product collection, and during the washing phase (B2), it penalises product loss, while α and γ promote the removal of impurities. Further to that, the constant δ introduces a small time penalty proportional to the duration spent in each phase, to discourage excessively long operation within any given step. While these weights are selected empirically, their relative magnitudes are grounded in domain-specific knowledge and were refined through targeted exploratory testing during preliminary training to ensure stable learning behaviour and physically meaningful control actions.

Table 7.2: Reward function weights.

Weights	**Phase I1**	**Phase B1**	**Phase I2**	**Phase B2**
Weak impurities (α)	10	0	0	5
Product (β)	5	10	5	10
Strong impurities (γ)	0	0	10	10
Time penalty (δ)	0.05	0.05	0.05	0.05

In addition to the phase specific weights, the magnitude of the flowrate (Q_t) directly scales the reward in all phases. As such, the agent is not only incentivised to identify which phase-specific behaviours yield high rewards, but also to modulate the flowrate accordingly. For example, lower values of Q_t may be advantageous in recycling phases (I1 and I2) to minimise product loss. In contrast, higher values may be preferred during collection (B1) or impurity removal (B2) to maximise throughput and efficiency.

This modular reward structure allows the controller to learn distinct flowrate behaviours appropriate to each process phase. Importantly, it also shapes the agent's implicit understanding of when transitions between phases should occur. Because the agent is not directly responsible for initiating phase transitions, these are implemented procedurally within the environment based on observed

concentration trends.

More specifically, during phase I1, the transition to B1 occurs when the outlet purity reaches or exceeds the specified threshold, or when the product concentration begins to decline after reaching a local maximum, which signals the peak of the product band. In phase B1, product is collected until purity falls below the threshold, at which point the controller transitions to phase I2. This ensures that product is collected only when it is within the required purity specifications. The recycling phase I2 then continues until the product concentration falls below a negligible threshold, indicating the end of the elution tail and prompting a switch to the B2 washing phase. Finally, the transition between phases B2 and I1 occurs when the product first breaks through the left-hand column, at which point, the columns switch roles (Figure 3.1). The controller is thus trained within a structured sequence of operations, but learns to regulate flowrates in ways that complement these transitions and enhance phase efficiency.

Finally, to ensure long-term efficiency, an additional penalty is applied at the end of each process cycle. This penalty scales with the number of timesteps required to complete the cycle, thus discouraging unnecessary extension of steps and promoting timely transitions. The penalty is computed based on the number of steps required to complete the preceding cycle and is then added as a one-off negative reward at the beginning of the new cycle, serving as a feedback signal for flowrate optimisation across the full process cycles.

7.2.5 Controller training

The RL controller is trained using the high-fidelity mechanistic model of the twin-column MCSGP system as a simulated environment. Each training episode consists of five process cycles and begins with randomised feed concentrations and a randomly selected product purity threshold. This variability forces the agent to learn adaptable control strategies rather than overfitting to a fixed scenario, enhancing robustness across different process conditions.

Training is carried out using the PPO algorithm, a policy-gradient method that supports stable learning under partial observability and continuous control. The environment provides stepwise rewards shaped by phase-specific objectives, while also applying cycle-level penalties to discourage unnecessarily long operations. Episodes terminate after five completed cycles, balancing short-term phase

optimisation with long-term performance efficiency.

At the core of the RL implementation is a feedforward actor-critic network (Figure 7.2). This architecture consists of a shared backbone, which comprises two hidden layers with 64 neurons each and employs the ReLU activation function, feeding into two separate output layers. The actor maps observations to flowrate actions, while the critic estimates the value function, representing the expected cumulative future reward from each state. This estimate supports the actor's learning by providing a baseline for evaluating the advantage of each action. Sharing the base layers of the network enables better representation learning, as both the policy and value function can benefit from common underlying features extracted from the input space.

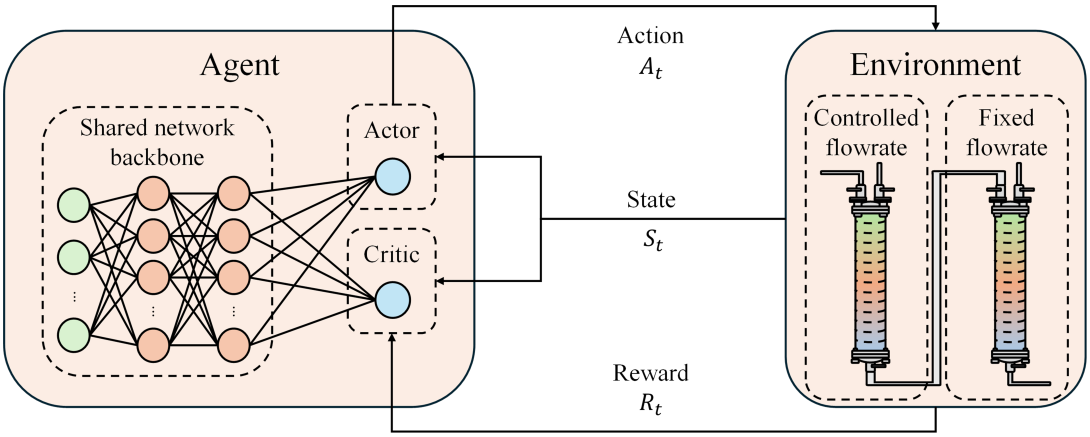


Figure 7.2: Reinforcement learning agent actor-critic network structure.

The policy outputs the mean of a Gaussian distribution over actions, enabling stochastic exploration during training. Actions are sampled, passed through a tanh transformation to bound them within $[-1, 1]$, and then rescaled to match the physical flowrate range. This avoids the discontinuities introduced by hard clipping and supports smoother, physically meaningful actions. A linearly decaying standard deviation balances exploration and exploitation, encouraging broader sampling in early episodes and more deterministic behaviour as the training progresses.

Episodes are seeded randomly to ensure broad state space exploration and improve generalisability. The initial training is conducted over 100 episodes, each beginning at process start-up and comprising five full process cycles, resulting in a total of 500 simulated cycles for policy learning. Once training is complete, the best-performing policy is stored and used for evaluation.

7.3 Results & discussion

7.3.1 Controller performance

The primary objective of the controller is to ensure that the purity specification is satisfied from start-up, while yield serves as a secondary measure of process efficiency. As described in the methodology, the policy is trained with the purity threshold as a supervisory input, allowing the product collection window to adapt dynamically during operation. The results below therefore focus first on the system trajectory from start-up to CSS, examining how quickly stable cyclic operation is established, how the controller regulates flowrates, and how the elution profiles of product and impurities evolve across cycles. Performance is then compared against nominal fixed-schedule operation and evaluated under different purity thresholds.

7.3.1.1 Nominal conditions

The performance of the RL controller is first evaluated under nominal operating conditions using a purity threshold of 98%. Figure 7.3 shows the outlet concentration profiles of the weak impurities, product, and strong impurities eluted from the right-hand column of the system at each phase (Figure 3.1) across the first ten process cycles, together with the corresponding flowrate trajectory applied by the controller.

The outlet concentration profiles in Figure 7.3a show that the system progressively approaches a cyclic profile, with the product and impurity peaks gradually stabilising across the first ten cycles. In the initial cycles, the species' elution profiles show significant differences in terms of peak intensity from cycle-to-cycle. As the operation progresses, the profiles evolve smoothly until, by cycle 10, the elution curves of consecutive cycles are indistinguishable, confirming that CSS has been reached. The attainment of CSS was verified using the following cycle-to-cycle convergence criterion (Equation 3.13). A tolerance of $\varepsilon_{\text{tol}} = 10^{-5}$ was adopted, under which the RL-controlled process reached CSS in ten cycles, similar to the nominal operating case (Papathanasiou et al., 2016). This indicates that the controller is able to reduce the transient length, thereby minimising the number of non-productive cycles that would typically need to be discarded in conventional operation (Mazzotti et al., 1997).

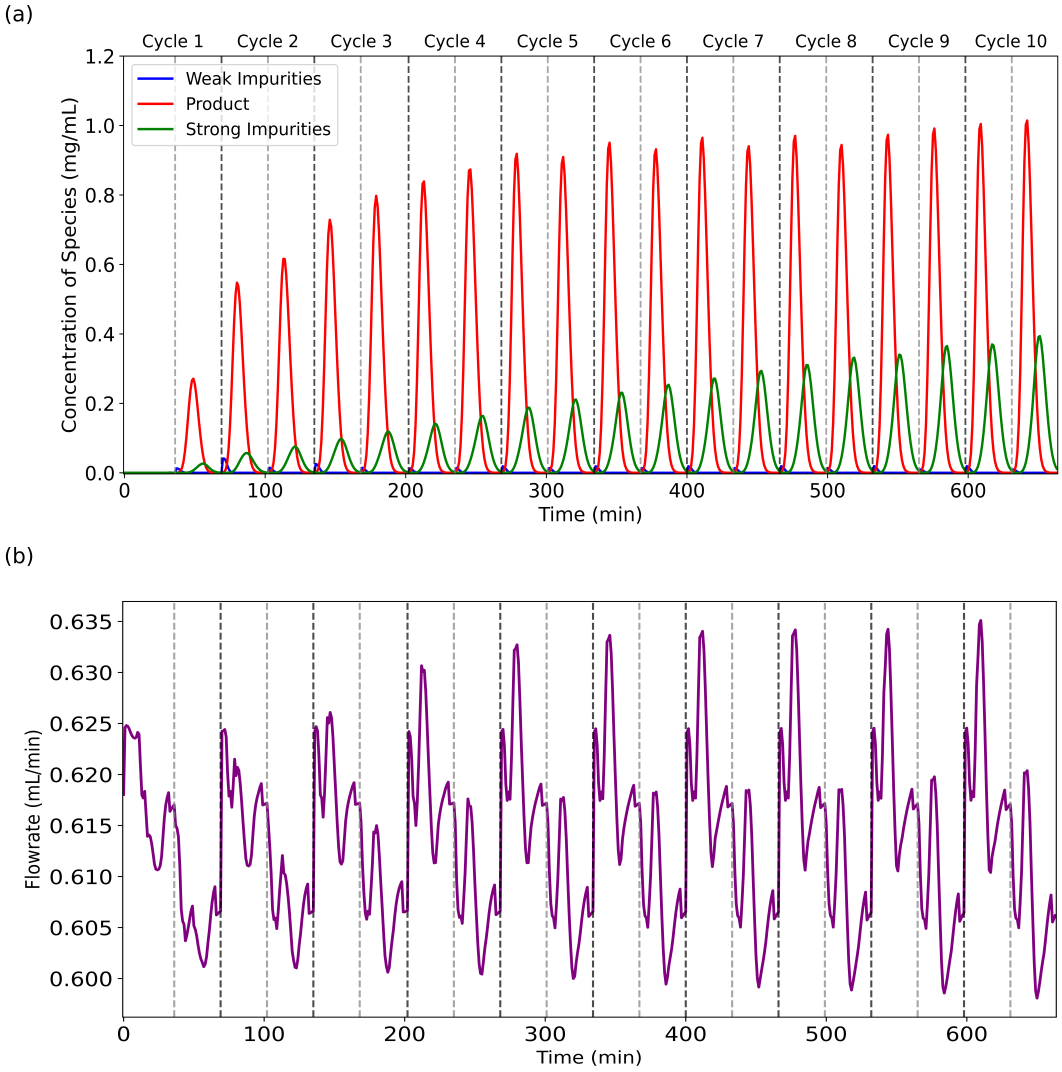


Figure 7.3: (a) Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process, the column responsible for the collection of the product and the recycling of the impure fractions, from process start-up across ten process cycles under RL control. (b) Flowrate applied by the controller on the right-hand column of the twin-column MCSGP system across the same cycles. The grey dashed lines represent the end of the first half of each process cycle, while the black dashed lines represent the completion of the second half of the cycle.

The corresponding flowrate trajectory applied by the controller is shown in Figure 7.3b. The profile reveals that the agent consistently selects values within a relatively narrow band around the middle of the allowable range. This behaviour suggests that the policy regulates the separation conservatively, avoiding extremes of very high or very low flowrate. Operating in this central region provides stability in the elution dynamics, allowing the controller sufficient resolution to identify when phase transitions should occur while ensuring that the product purity specification is maintained.

The flowrate trajectory also illustrates how the policy interprets the cyclic nature of the process. Despite not being provided with a pre-defined switching schedule, the controller establishes a periodic flowrate pattern from the outset. The oscillatory structure of the trajectory is synchronised with the cycle length and is reproduced consistently across cycles, confirming that the policy has internalised the underlying periodicity of the process. It is worth noting that, although the pattern of flowrates is repeated from one full cycle to the next, it is not identical between the two half-cycles of a given cycle. This results in the two columns converging to slightly different, but functionally equivalent, CSS profiles, as can be observed from cycle 10 in Figure 7.3a. Despite these minor differences, both columns operate with identical phase durations and yield nearly indistinguishable outlet concentrations, confirming that the controller enforces the cyclic symmetry of the process, without need for specific cycle information.

A further observation from the outlet concentration profiles (Figure 7.3a) is the steady accumulation of strong impurities in the system. This behaviour arises from their co-elution with the product during the recycling phases. Due to their stronger interaction with both the stationary phase and the product, strong impurities are inherently more difficult to separate completely from the system. Their presence contrasts with that of the weak impurities, which are efficiently removed during the early phases and are not recycled to the same extent. The RL controller accommodates this behaviour by maintaining flowrates that balance product recovery with impurity removal, preventing product loss while allowing gradual stabilisation of the impurity profiles across cycles.

Finally, it is important to note that the purity specification is satisfied from the first product collection window. All collected product fractions exceeded the 98% threshold, demonstrating that the controller did not require a stabilisation period to enforce the target purity. This distinguishes the RL-controlled opera-

tion from conventional nominal strategies, in which product is only guaranteed to meet purity specifications after the attainment of CSS (Rajendran et al., 2009). Ensuring compliance from start-up provides a significant advantage, as it allows the immediate collection of on-spec product without discarding early cycles.

7.3.1.2 Performance at CSS

To further analyse the controller-imposed operating profile, a single cycle at CSS under RL control is compared directly against the nominal CSS operation. Figure 7.4 presents the outlet concentration profiles of the weak impurities, product, and strong impurities for both cases, along with the corresponding product collection windows.

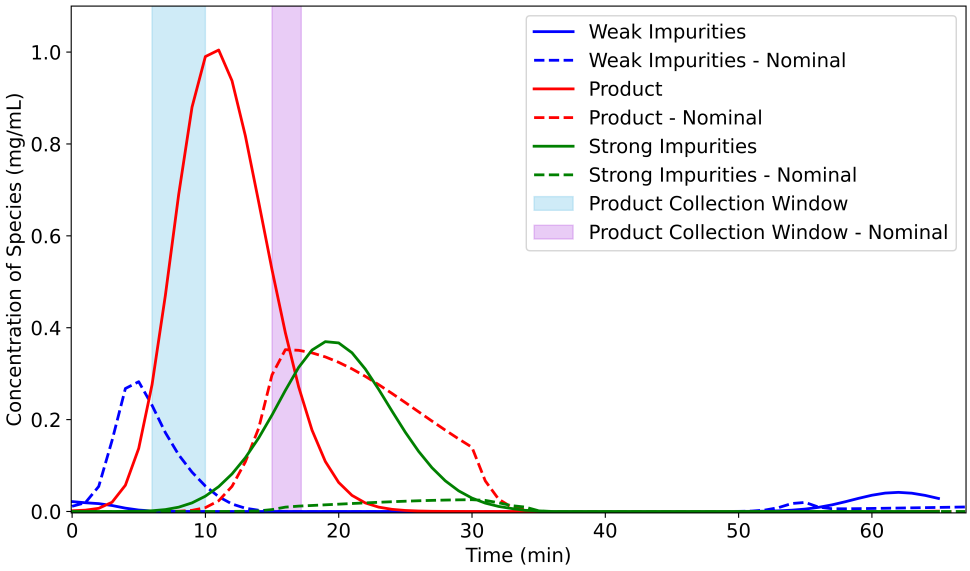


Figure 7.4: Comparison of the outlet concentrations of the weak impurities (blue), product (red), and strong impurities (green) for one process cycle at CSS, shown for the RL-controlled process profile (solid lines) and the nominal operating profile (dashed lines).

The comparison highlights several key differences between the two operating strategies. Firstly, the RL-controlled cycle is slightly shorter than the nominal one, even though no explicit constraint on cycle time was imposed during training. The only timing-related term in the reward function was the small penalty on phase duration. The fact that the controller converges to a shorter cycle length therefore indicates that the policy autonomously identified a more efficient balance between phase duration and overall reward, reducing operating time without external enforcement.

Within this reduced cycle duration, however, the distribution of the outlet profiles of the species differs significantly, both in the timing of elution and in the shape and height of the peaks. The RL-imposed product collection window is broader than in the nominal case and coincides with the period during which fresh feed is introduced into the system. As such, a larger amount of material is processed within each cycle. This results in a higher intensity of the product peak within the collection window and consequently a greater quantity of product recovered, as calculated from the applied flowrates, the corresponding product concentrations, and the duration of the collection window. Table 7.3 shows that the RL controller achieves almost double the absolute product recovery compared to the nominal operation, while maintaining a comparable yield and purity. This demonstrates that the policy reallocates the cycle in a way that preserves separation efficiency but extracts greater value from the available operating time, by increasing the total amount of product treated and therefore purified per cycle.

Table 7.3: Comparison of performance metrics for the RL-controlled and the nominal operating profile for one process cycle at CSS.

Operation	**Purity (%)**	**Yield (%)**	**Product recovery (mg)**
Nominal	98	86	1.40
RL-controlled	98	85	2.52

The differences between the two strategies are not limited to the product elution profile. The weak impurities, which elute ahead of the product peak, show reduced overlap with the product elution window under RL control compared to the nominal case, further supporting efficient product collection. By contrast, the strong impurities display greater accumulation in the RL profile. This behaviour is consistent with the observations made during the ten-cycle analysis, but here the underlying cause can be more clearly identified. In the reward function, phase I2, which is responsible for recycling the overlapping product and strong impurities fraction, was designed to positively reward recycling of mixed streams. This structure ensured that the controller did not recycle product-rich fractions that already met the purity specification and further encouraged the separation of the product from the impurities. However, the small positive weight assigned to the impurities concentration in this phase also permitted gradual co-recycling of

strong impurities alongside the intended product, resulting in the accumulation observed here.

Taken together, these results confirm that the RL policy does not alter the nominal timing of the cycle but instead reallocates performance within it, increasing throughput while maintaining purity and yield. This illustrates the controller's ability to extract greater value from the same operating window, highlighting the benefit of the adaptive strategy compared to pre-optimised nominal schedules.

7.3.1.3 Purity threshold

To demonstrate the flexibility of the RL controller, its performance is also evaluated under nominal conditions with a reduced purity threshold of 95%. Figure 7.5 compares the elution profiles at CSS for the two thresholds, along with the corresponding flowrate trajectories applied by the controller. The overall shapes of the profiles remain highly consistent, and the onset of product elution occurs at the same time in both cases, showing that the controller does not alter the fundamental timing of the separation. However, reducing the purity requirement results in a broader product collection window, as the purity constraint is satisfied for a longer period. This adjustment reduces the maximum concentration of the product peak but also leads to less recycling of material and therefore a smaller accumulation of impurities across cycles.

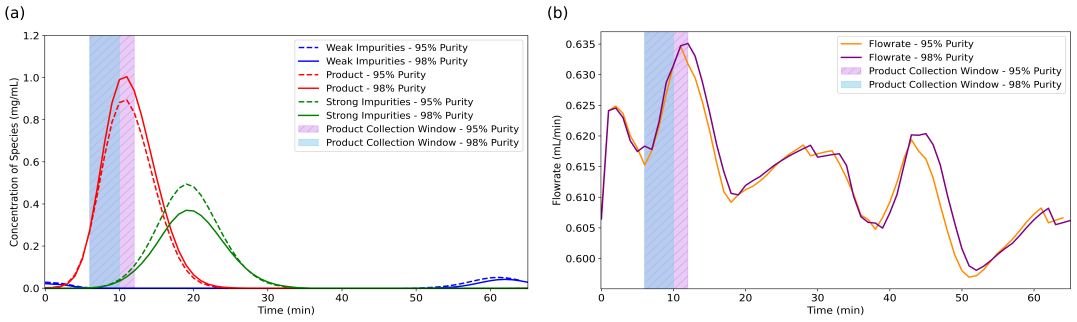


Figure 7.5: Comparison of (a) outlet concentration profiles of the weak impurities (blue), product (red), and strong impurities (green), and respective (b) flowrate trajectories at CSS for RL-controlled operation under 98% (solid lines) and 95% (dashed lines) purity thresholds.

The differences between the two thresholds are more clearly seen in the phase durations reported in Table 7.4. Although no constraint on cycle length was

imposed during training, the controller converges to the same overall cycle duration in both cases, which suggests that it autonomously identifies a stable timing structure for the process. The main differences appear in the allocation of time within the cycle. At the 95% purity threshold, the collection phase B1 is extended by two minutes relative to the 98% case, while the subsequent recycling phase I2 is shortened by the same amount. With a broader acceptable purity range, the controller allows direct capture of a greater portion of the product peak and correspondingly reduces the time spent recycling overlapping fractions.

Table 7.4: Phase durations at CSS for RL-controlled operation under 98% and 95% purity thresholds.

	I1 (min)	**B1 (min)**	**I2 (min)**	**B2 (min)**
98% threshold	5	4	20	4
95% threshold	5	6	18	4

These results show that lowering the purity threshold does not alter the overall separation pattern but shifts the balance between collection and recycling. The RL controller adapts to the new specification without retraining, broadening the product collection window and reducing recycling demands, while converging to the same cycle duration. This behaviour highlights the controller's ability to flexibly adjust to different supervisory inputs, ensuring compliance with variable quality requirements while continuing to operate efficiently.

7.3.2 Performance under feed variability

To further evaluate the robustness of the RL controller, the system was subjected to variability in the feed composition. Feed composition is considered a measured disturbance, since the characteristics of a given batch can be determined in advance and, in continuous chromatography, monitored during operation using PAT tools (Rüdt et al., 2017). Two representative scenarios were examined. In the first, the feed composition was made more favourable by increasing the product concentration by 20% while decreasing the concentrations of both weak and strong impurities by 20%. In the second, a more challenging condition was introduced by reducing the product concentration by 20% while increasing both impurities by 20%. These scenarios reflect realistic fluctuations in upstream titre

and impurity levels, while remaining within the validation range of the mechanistic model.

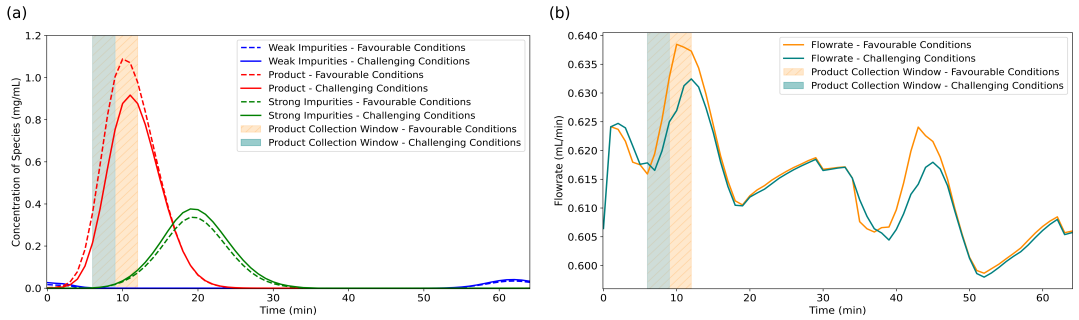


Figure 7.6: Comparison of (a) outlet concentration profiles of the weak impurities (blue), product (red), and strong impurities (green), and respective (b) flowrate trajectories for RL-controlled operation under two feed scenarios at CSS. Favourable conditions (solid lines) correspond to a 20% increase in product concentration and 20% decrease in impurities, while challenging conditions (dashed lines) correspond to the opposite change.

The outlet concentration profiles and corresponding flowrate trajectories for the two cases are shown in Figure 7.6. Under favourable conditions, the improved product-to-impurities ratio allowed the controller to broaden the product collection window, which coincided with a sharper and more pronounced product peak. The flowrate trajectory in this case exhibited distinct spikes within the collection phase, steepening the product front and enabling direct recovery of a larger fraction of the elution peak. By contrast, under the challenging feed composition, where the overlap between product and impurities was more intense, the collection window narrowed to half of its duration compared to the favourable conditions case. The flowrate trajectory was also more restrained, with fewer and smaller spikes, indicating that the policy maintained purity by limiting the time spent in collection and relying more heavily on recycling to separate the overlapping fractions.

Across both scenarios, the purity specification was consistently enforced, with the controller reallocating time between collection and recycling in line with the feed composition. This resulted in a broader and more aggressive collection under favourable conditions, where recovery could be increased without compromising quality, and a shorter and more conservative collection under challenging conditions, where purity was preserved at the expense of immediate recovery. These behaviours align with the structure of the reward function and the phase-

dependent switching logic, and they demonstrate the ability of the RL controller to adapt its decisions to measured disturbances in feed composition.

7.3.3 Performance against disturbances

In addition to measured variability in the feed composition, the robustness of the controller was also tested against disturbances in the applied flowrate while operating at CSS. Such disturbances may arise from pump malfunctions or valve switching errors and cannot be anticipated with the same fidelity (Krättli et al., 2013b). They therefore represent a more challenging class of disturbance, as the controller must react solely on the basis of the observed outlet signals.

A set of disturbance scenarios was considered comprising flowrate disturbances lasting one minute at different phases of the process after reaching CSS. The disturbances were implemented in the right-hand column of the system, at timepoints when product elution is in progress. To avoid bias in their placement, the timepoints within the phases of the cycle that the disturbances occurred were selected at random, corresponding to the fourth timestep of phase I1, the second timestep of phase B1, and the eighth timestep of phase I2. Figure 7.7 shows the imposed flowrate profiles containing the disturbances together with the resulting outlet concentrations over three consecutive cycles. In each case the disturbance was implemented by forcing the flowrate of the right-hand column to its maximum value for the specified duration.

As shown in Figure 7.7, all three disturbance scenarios lead to similar changes in the outlet concentration profiles. In each case, the product peak in the cycles following the disturbance is reduced in height, while the weak and strong impurities remain mostly unaffected. These deviations reflect the fact that a temporary increase in flowrate accelerates elution and shortens the duration of the product collection phase B1, as can be observed in the first two disturbance scenarios (Figure 7.7a,c), which also governs the introduction of fresh feed into the system. With less material entering, the subsequent cycle operates with a reduced mass balance, which is visible in the lower product peaks. A similar effect is observed in the third scenario, where the disturbance takes place in phase I2, after the product collection window. In this case the separation of the product from the impurities is affected. This is then reflected in the second half of the cycle, where the duration of the recycling phase I1 in the other column of the

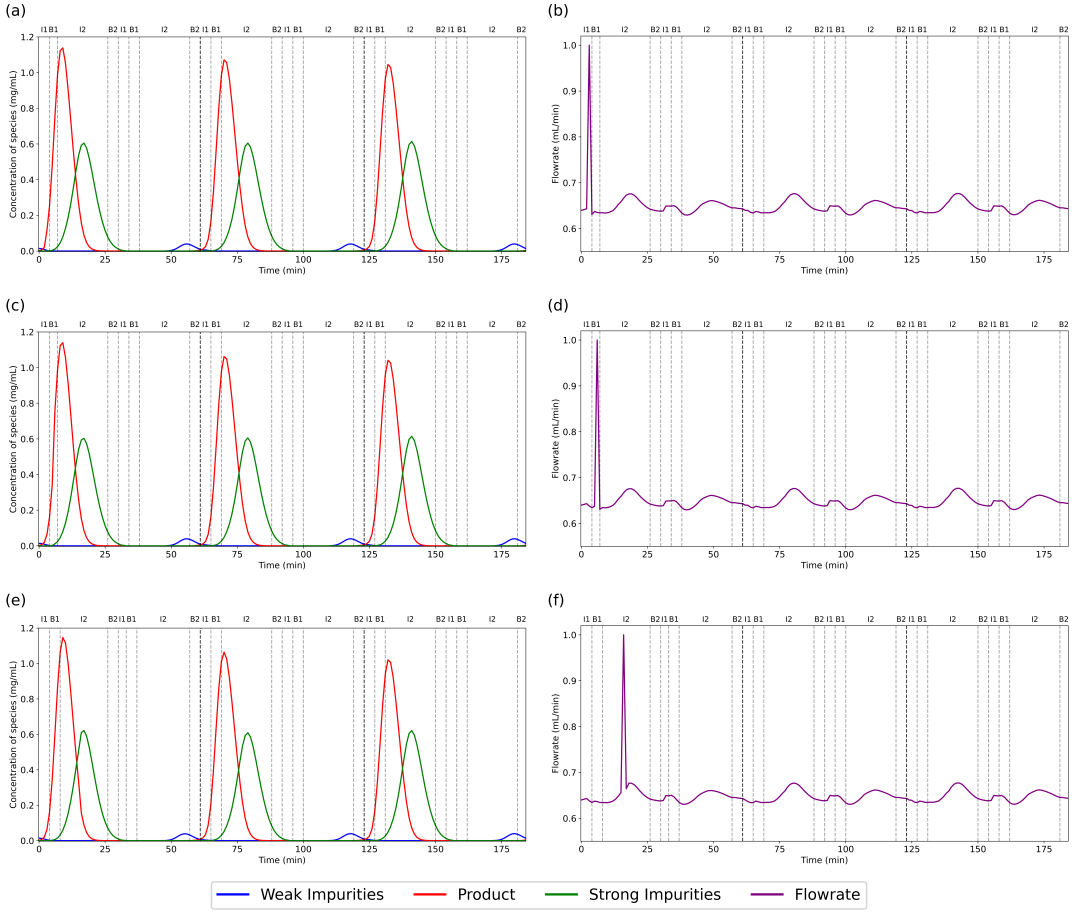


Figure 7.7: Outlet concentration profiles of weak impurities (blue), product (red), and strong impurities (green) together with the applied flowrate trajectory (purple) for three process cycles under RL control at CSS in the presence of one-minute disturbances. The disturbance was applied at the fourth timestep of I1 (a,b), the second timestep of B1 (c,d), and the eighth timestep of I2 (e,f). The grey dashed lines represent phase switches and the black dashed lines represent the completion of one cycle.

system is reduced (Figure 7.7e). As such, less product is recycled and therefore the elution peak observed in the cycle after the disturbance is reduced.

Importantly, the degree of perturbation in the elution profiles is comparable across phases I1, B1, and I2, as are the flowrate trajectories applied by the controller after the disturbances have occurred (Figure 7.7b,d,f). This indicates that short disturbances do not have a lasting effect on the process across many cycles and are absorbed by the controller. In addition, the controller maintains the product purity above the specified threshold in all cases, despite the disturbances. This is achieved by adjusting the switching times to protect the collection window, ensuring that only on-spec material is captured. Although the trade-off

is a temporary reduction in product throughput due to the lower feed load, the system converges back to CSS within three cycles, confirming that the controller can effectively manage short disturbances while maintaining product quality.

To further test the controller against possible deviations from normal operation, a second set of disturbance scenarios was considered, is which the duration of the previously applied flowrate disturbances was extended to three minutes. Figure 7.8 shows the resulting outlet concentration profiles and flowrate trajectories over three cycles following the disturbance.

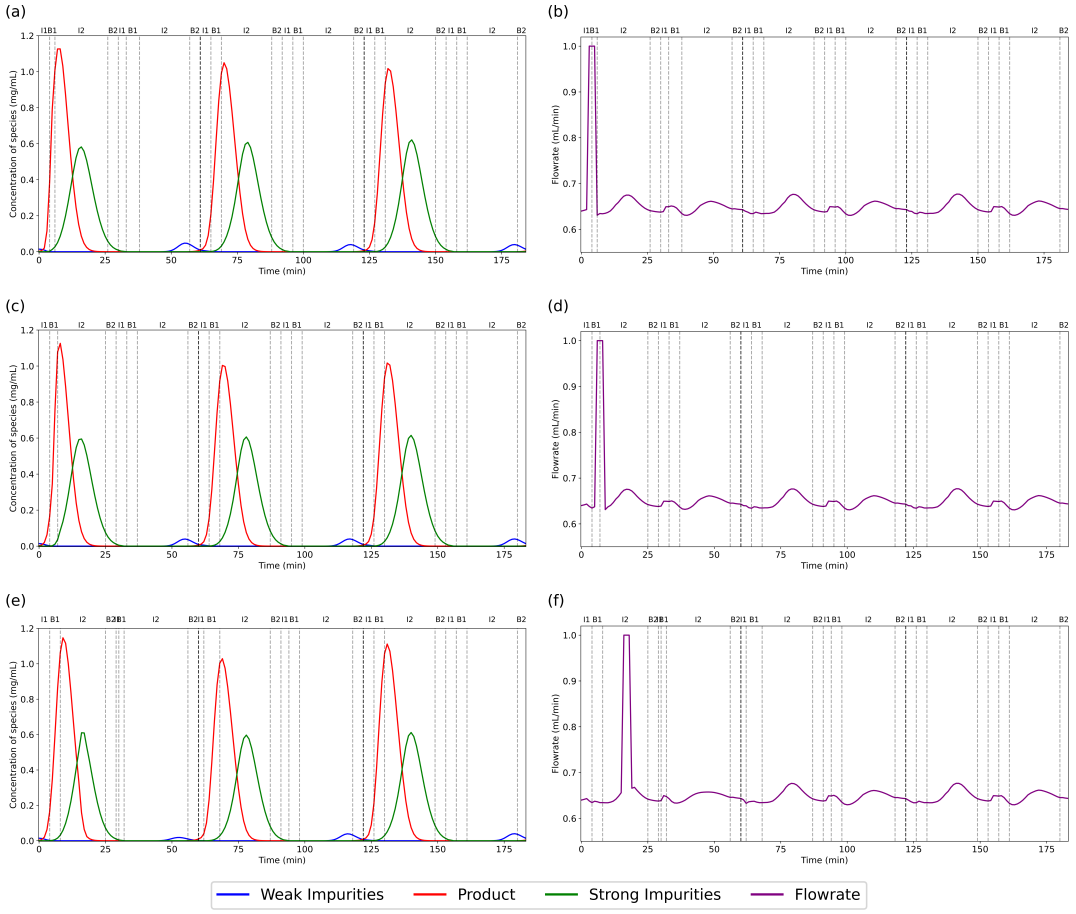


Figure 7.8: Outlet concentration profiles of weak impurities (blue), product (red), and strong impurities (green) together with the applied flowrate trajectory (purple) for three process cycles under RL control at CSS in the presence of three-minute disturbances. The disturbance was applied at the fourth timestep of I1 (a,b), the second timestep of B1 (c,d), and the eighth timestep of I2 (e,f). The grey dashed lines represent phase switches and the black dashed lines represent the completion of one cycle.

Based on Figure 7.8, the effect of the longer disturbance on the process operation was more pronounced, particularly when it overlapped with the product

collection phase B1. When the disturbance occurred in phase I1, the outlet profiles of the weak and strong impurities, as well as the product, were impacted in terms of peak height and width (Figure 7.8a). The disturbance expedited the transition to B1 and shortened the subsequent collection phase, leading to a narrower and taller product peak compared to CSS conditions. Because B1 was shortened, less fresh feed was introduced into the left-hand column, which was reflected in reduced product peaks in the following cycles. The controller partially compensated within the same cycle by extending the B1 phase in the second half, thereby attempting to counterbalance the reduced mass introduced earlier. In the next two cycles the phase durations and flowrate trajectory progressively realigned with those at CSS, indicating that the policy was steering the system back towards CSS (Figure 7.8b).

Similarly, the disturbance in phase B1 caused the product collection window to be shortened, sharpening the product peak and accelerating the onset of recycling (Figure 7.8c,d). Although the purity specification was preserved, the amount of product collected in the disturbed cycle was reduced, since a larger fraction was diverted to recycling. This was accompanied by an increase in the strong impurities peak in the following cycles, consistent with the recycling of more material in the system, as well as the reduced separation from the product due to the reduced residence time. Nevertheless, the controller was able to start guiding the system back towards CSS from the second cycle onwards, by applying the flowrate trajectory and phase durations employed in CSS.

In the case of phase I2, the disturbance primarily affected the recycling of the overlapping product and strong impurities fraction (Figure 7.8e,f). The immediate effect was a sharper peak in the strong impurities elution profile in the disturbed cycle, accompanied by a very short collection window in the second half of the cycle, during which the product only satisfied the purity threshold for a brief period. The reduced duration of phase B1 meant that less feed entered the system, which was reflected in the lower peaks of all species in the next cycle. The disturbance also carried over to the first half of the second cycle, where the duration of the product collection phase was increased compared to CSS operation, as the co-elution with weak impurities was shortened, thus increasing the fresh feed introduced in the system. By the third cycle, however, both the phase durations and the flowrate trajectory were restored to their nominal values, showing that the controller was able to begin re-establishing CSS despite

the extended perturbation.

Across all three disturbance scenarios, the controller maintained the product purity above the specified threshold by adapting the phase switching times, without increasing the duration of the process cycles. When the one-minute and three-minute cases are compared directly, several differences emerge. In phases I1 and B1, short disturbances were absorbed almost immediately, with only minor perturbations to the outlet profiles, whereas the longer disturbances produced more visible deviations in the product peaks and reduced feed introduction in the disturbed cycle. In both cases the controller restored CSS, but the three-minute disturbances required additional corrective action, such as extending phase B1 in subsequent half-cycles, before full recovery was achieved. Interestingly, disturbances in phase I2 proved more disruptive, as they directly interfered with the recycling and more importantly the effective separation of overlapping product and strong impurities fractions rather than simply accelerating the elution of already pure product. This distinction meant that while I1 and B1 phase disturbances shared similar profiles and were compensated quickly, the I2 phase disturbance had a more lasting impact on the separation, carrying into subsequent cycles and requiring additional corrective actions by the controller. Across all scenarios, the purity specification was maintained, though at the cost of lower product peaks. Taken together, these results show that while the RL controller enforces product quality under both short and extended disturbances, the placement and duration of the disturbance governs the magnitude of the deviation and the speed of recovery, with longer faults amplifying the trade-off between purity satisfaction, productivity, and cycle efficiency.

7.3.4 Measurement noise

The robustness of the controller was further evaluated by introducing measurement noise into the observable signals. Noise up to 5% was applied on the pooled outlet concentration, which represents the signal available from the chromatogram (Figure 7.9). This simulates realistic operation, where detector variability is unavoidable. To preserve consistency between components, the perturbed total concentration was redistributed among the weak impurities, product, and strong impurities according to their relative contributions as defined by the mechanistic model. In this way, the model was used as an auxiliary tool,

ensuring that the noise was applied in a physically meaningful manner and that the component profiles remained representative of the separation taking place.

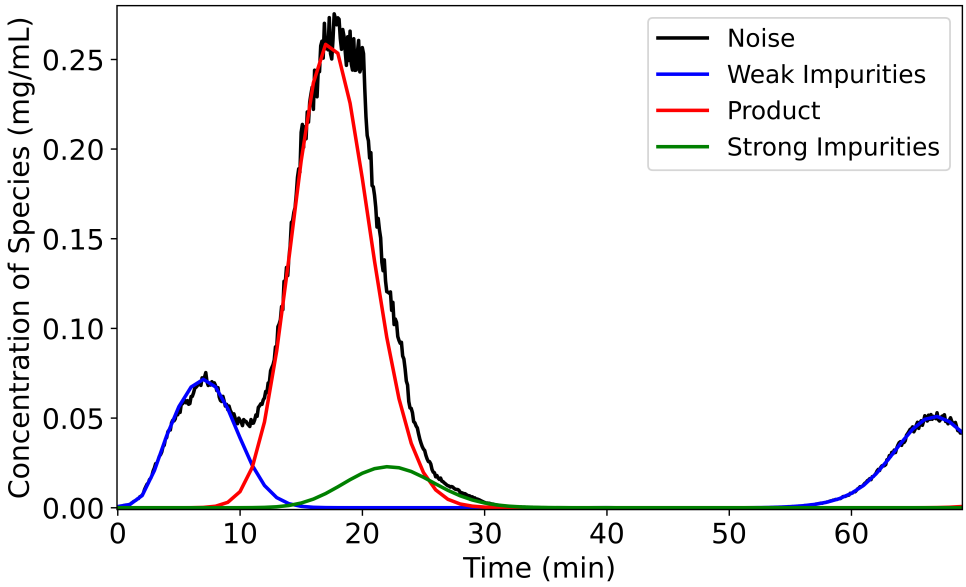


Figure 7.9: Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process for one process cycle under RL control, in the presence of measurement noise up to 5% applied to the pooled outlet signal (black).

Figure 7.10 shows the outlet concentration profiles and flowrate trajectories across ten cycles from process start-up. The overall elution behaviour closely matches the noise-free case (Figure 7.3). The effect of the added noise is visible as small fluctuations on the product and impurities peaks, but these do not alter the timing of elution or the separation between species. Similarly, the flowrate trajectory shows minor adjustments in the magnitude of selected flowrate values, reflecting the controller's response to noisy signals, yet the overall trajectory and phase structure remain consistent with nominal operation.

Despite the imposed variability, the controller maintained the purity specification throughout all cycles. Phase durations were stable, and switching decisions continued to align with the underlying elution dynamics. The trade-off introduced by noise was therefore limited to small fluctuations in observed concentrations and corresponding minor adjustments in flowrate levels, with no impact on the overall separation pattern or performance. These results confirm that the controller is robust to realistic levels of measurement noise, preserving stable operation and enforcing product quality even under noisy measurements.

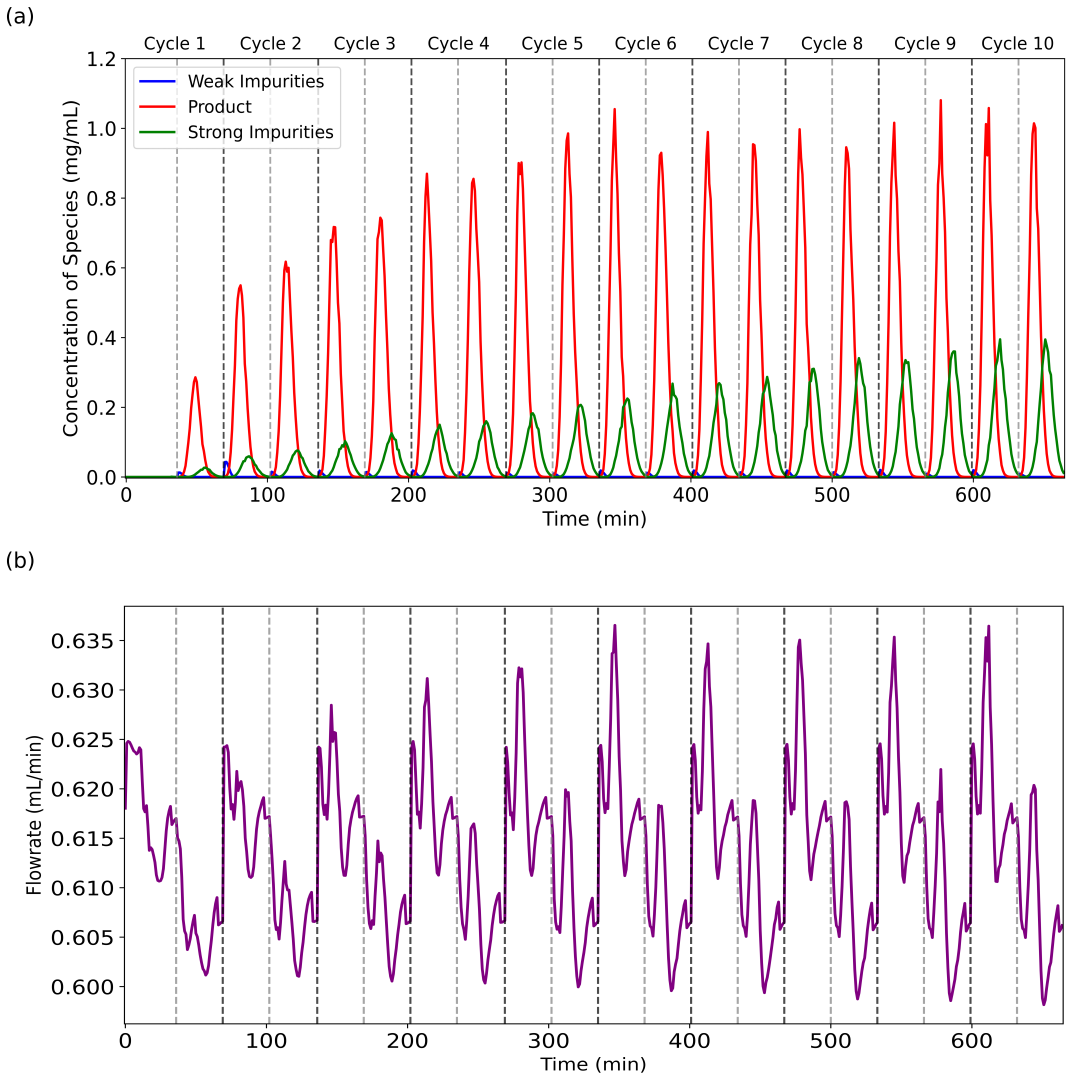


Figure 7.10: (a) Concentrations of the weak impurities (blue), product (red), and strong impurities (green) being eluted from the right-hand column of the twin-column MCSGP process, the column responsible for the collection of the product and the recycling of the impure fractions, from process start-up across ten process cycles under RL control, in the presence of measurement noise of up to 5%. (b) Flowrate applied by the controller on the right-hand column of the twin-column MCSGP system across the same cycles. The grey dashed lines represent the end of the first half of each process cycle, while the black dashed lines represent the completion of the second half of the cycle.

7.4 Conclusions

In this Chapter, an RL controller was developed and evaluated for the real-time operation of the twin-column MCSGP process. The controller was trained based on a high-fidelity mechanistic model used as the process environment to learn adaptive flowrate and switching time regulation across all process phases and to adapt to different product purity requirements. The performance of the controller was assessed under nominal, disturbed, and noisy conditions to evaluate its robustness and flexibility.

Under nominal operation, the controller was shown to autonomously converge to stable cycle operation, enforcing purity requirements from start-up and adapting flexibly to different supervisory thresholds without the need for retraining. When subjected to feed variations and flowrate disturbances, it maintained the purity constraint while re-establishing CSS through adaptive phase timing, and under realistic levels of measurement noise it preserved stable operation, with only minor fluctuations in flowrate and outlet signals. Across all scenarios, the controller demonstrated robustness and reliability, consistently ensuring product quality while adapting its operation to perturbed conditions.

Placed in the wider context of this thesis, these results highlight how RL complements the modelling and optimisation strategies developed in the earlier Chapters. While data-driven and hybrid models enable efficient online optimisation, RL provides a framework for adaptive, phase-aware control that does not rely on predefined switching rules. Together, these tools illustrate a cohesive approach to the digitalisation of chromatographic separation, where modelling, optimisation, and control act in synergy to enhance efficiency, resilience, and real-time applicability.

Chapter 8

Conclusions & outlook

8.1 Summary of contributions

This thesis presents the development of machine learning (ML) frameworks for modelling, optimisation, and control of continuous chromatographic separation processes. Current industrial practice highlights the importance of rapid and reliable tools to support process development and operation, yet conventional mechanistic models can be computationally prohibitive for real-time use. In this context, surrogate modelling and reinforcement learning (RL) provide the opportunity to deliver predictive accuracy at a fraction of the computational cost, enabling optimisation and adaptive control under the constraints of continuous biomanufacturing. Taken together, these developments form a modelling toolbox that extends the role of mechanistic models into domains where computational demands would otherwise preclude their application.

The review of literature presented in Chapter 2 showed that mechanistic models dominate chromatographic process development, but their complexity often makes them constrained in online applications, such as optimisation and closed-loop control. Data-driven models have been increasingly adopted to overcome this barrier, providing rapid evaluation and the potential for integration with real-time decision-making, though their performance is often restricted to the operating domains represented in the training data. Hybrid formulations have emerged more recently as a promising compromise, combining mechanistic structure with data-driven flexibility, but their systematic application to chromatography remains limited. In further applications, such as optimisation and control, most studies have been confined to offline investigations, with model-based approaches hindered by evaluation times that are too long to accommodate variability during operation. Together, these gaps motivated the development of surrogate and adaptive methods able to balance fidelity and speed within the strict time horizons of online decision-making.

The twin-column multicolumn countercurrent solvent gradient purification (MCSGP) process was introduced in Chapter 3 as the case study through which the modelling, optimisation, and control frameworks of this thesis were developed and assessed. Its cyclic operation, non-linear dynamics, and stringent purity requirements exemplify the types of processes where conventional optimisation and control encounter significant challenges. The mathematical formulation and system analysis underlined the computational demands of simulating the system

and reaching cyclic steady state (CSS), illustrating the need for surrogate models and adaptive control strategies to make such processes tractable for online applications.

Chapter 4 examined the potential of artificial neural networks (ANNs) as purely data-driven surrogates trained on mechanistic simulations of the twin-column MCSGP process at CSS. The models mapped operating conditions and inlet concentrations to full outlet trajectories of the separation species over a cycle, without requiring explicit time encoding. Two formulations were compared, a single multiple-input/multiple-output (MIMO) network and a set of independent multiple-input/single-output (MISO) networks, with their performance assessed against variations in architecture and training-set size. Both approaches reproduced elution behaviour with high accuracy under nominal and interpolated conditions, while reducing simulation time by three orders of magnitude. These results demonstrated that data-driven models can reproduce mechanistic behaviour efficiently within their training domain, while eliminating the computational burden associated with mechanistic models.

To improve the predicting capabilities of the developed data-driven models, Chapter 5 developed a hybrid model that integrated mechanistic knowledge, through the separation isotherm, with an ANN, eliminating the need for spatial discretisation. The model relied only on experimentally measurable quantities and achieved a 97% reduction in simulation time compared to the mechanistic model of the system. Validation confirmed that the hybrid maintained predictive accuracy under interpolation and showed reliability when applied to operating conditions outside the training domain. The results demonstrated that embedding mechanistic balances within a data-driven framework improves predictive performance while preserving computational efficiency, establishing the hybrid model as an alternative for applications where extrapolation and interpretability are essential.

Having established fast and accurate surrogate models, Chapter 6 integrated them into Bayesian optimisation (BO) workflows to evaluate their role in process decision-making. Both offline and online scenarios were considered. Offline optimisation confirmed that the data-driven and hybrid models converged to solutions consistent with the mechanistic model, ensuring product purity and demonstrating negligible differences in the achieved process yield. Under online horizons, however, the mechanistic model failed to converge within the available

time, while both surrogate models maintained reliable performance. The hybrid model delivered the closest match to the mechanistic solutions, while the purely data-driven model converged more rapidly with minor trade-offs in yield. These results demonstrated the potential of surrogate models to enable real-time optimisation in chromatographic processes.

Finally, Chapter 7 expanded the developed modelling toolbox beyond optimisation by introducing RL for adaptive control. A Proximal Policy Optimisation (PPO) agent was trained on mechanistic simulations to regulate the flowrate and switching times of the process with the objective of maintaining product purity. The agent observed process states including species concentrations, operating step, and flowrate, and learned policies that adapted dynamically over full cycles, catering to variable product purity specifications. Once trained, the controller proved robust across a range of scenarios and responded effectively to variability in feed composition, process disturbances, and measurement noise. The results showed that RL can provide stable operation from start-up and achieve adaptive control under conditions, thereby enhancing process robustness.

Next, the main conclusions that result from the findings of this thesis are outlined (Section 8.2), followed by a discussion of future research directions that can build on this work (Section 8.3).

8.2 Main conclusions

The main conclusions that can be drawn from the studies presented in Chapters 4, 5, 6, and 7 are listed below:

8.2.1 Data-driven modelling

ANNs can serve as efficient surrogates for the twin-column MCSGP process, reducing evaluation times by 99% while maintaining accuracy within their training domain. Their reliability depends on network structure, with ReLU activated single-output formulations producing smoother trajectories and more consistent optimisation performance. These findings establish data-driven surrogates as viable tools for rapid evaluation, though their generalisability remains limited to validated operating spaces.

8.2.2 Hybrid modelling

Hybrid models that embed mechanistic knowledge extend predictive capability beyond purely data-driven approaches. By retaining the isotherm while removing spatial discretisation, they deliver accuracy across a broader operating space, offering reliable predictions outside their training domain, and reduce computational cost. This shows that hybrid models offer a tractable balance between efficiency and fidelity, making them suitable for applications where interpretability and robustness are critical.

8.2.3 Process optimisation

Surrogate models make real-time optimisation feasible in cases where mechanistic models are too computationally demanding. Within BO frameworks, they can perform process optimisation under offline conditions and sustain performance under online horizons where mechanistic models are limited by their computational complexity. Importantly, the suitability of each model depends on its structure and the available time horizon for decision-making, with hybrid models suited to cycle-level optimisation and data-driven models enabling minute-by-minute decisions.

8.2.4 Reinforcement learning for adaptive control

RL can deliver stable and adaptive control of the twin-column MCSGP process under demanding operating conditions, adjusting flowrates and switching between process phases in real-time. It is able to satisfy stringent purity requirements from start-up, adjust dynamically to changes in feed composition, and remain robust in the presence of process disturbances and measurement noise. RL offers a flexible control paradigm that is well suited to environments where variability and uncertainty challenge conventional model-based strategies.

8.3 Future directions

This thesis has established a modelling toolbox for continuous chromatographic processes, combining data-driven and hybrid models with RL to enable optimisation and adaptive control under computational and operational constraints.

At its current stage, the toolbox can be applied to reproduce mechanistic predictions efficiently, support optimisation across offline and online horizons, and maintain robust operation of the twin-column MCSGP process under variability. These capabilities mark a step towards practical decision-support for intensified downstream processing.

The following sections focus on (i) extending surrogate models through transfer learning for new products, resins, and operating regimes, (ii) developing receding-horizon formulations for BO, (iii) broadening RL to dual-column control, and (iv) applying the methods to other intensified separation processes.

8.3.1 Transfer learning

The surrogate models developed in Chapters 4 and 5 demonstrated that both data-driven and hybrid approaches can reproduce mechanistic predictions efficiently while reducing computational cost. Their applicability, however, was limited to a single case study of the twin-column MCSGP process.

A next step of this work is the use of transfer learning to extend the applicability of the developed data-driven and hybrid models. At present, the models were trained on in-silico data generated for a single separation problem, with their predictive performance evaluated within and beyond the training space. In practice, new products, column resins, or operating conditions would require additional training data, which can limit the deployability of surrogate approaches. Transfer learning offers a route to adapt existing surrogates to related systems with minimal retraining, leveraging shared representations of underlying dynamics. Such strategies would reduce the data and computational burden associated with developing models for each new case.

8.3.2 Receding-horizon optimisation

The optimisation studies performed in Chapter 6 demonstrated that both data-driven and hybrid models make dynamic optimisation feasible, sustaining reliable performance at cycle-to-cycle and even minute-by-minute horizons where mechanistic models could not converge within the available time. These results confirmed that surrogate models extend optimisation from purely offline studies into online decision-making, adapting operating points within the time constraints of process operation.

Building on this, future work could adopt a receding-horizon formulation, where BO is re-solved as new process information becomes available. Instead of fixing operating conditions once per horizon, decisions would be updated continuously, enabling dynamic responses to feed variability and process disturbances. This would integrate optimisation more closely with real-time operation and further establish surrogates as active enablers of adaptive process control.

8.3.3 Dual reinforcement learning control

The RL controller presented in Chapter 7 showed that flowrate control of the right-hand column, executing the elution of the species, can maintain product purity from start-up and remain stable under feed variability, process disturbances, and measurement noise. In the present formulation, the flowrate of the left-hand column, responsible for the feeding and separation tasks, was held at its maximum, which limits flexibility to redistribute flows between columns.

A natural extension of this work is to expand the controller to both columns, enabling simultaneous regulation of their flowrates and providing greater flexibility in adapting to process variability. This dual-controller setup would increase the degrees of freedom available to the agent and allow more efficient management of trade-offs between purity, yield, and efficiency, particularly during start-up and step transitions where coordinated adjustments between columns can improve overall stability and performance.

8.3.4 Other applications

While this thesis has focused on the twin-column MCSGP process as a representative case study, the developed frameworks are not limited to chromatographic systems. Many intensified separation processes, such as crystallisation and pressure swing adsorption (PSA), share the features of cyclic operation, multiple interacting units, and stringent purity or recovery requirements (Orehek et al., 2020; Ward and Pini, 2022). The modelling and RL methods developed could be transferred to these systems with minimal adaptation, providing the same advantages of reduced computational cost and enhanced adaptability under uncertainty. Extending the toolbox to such applications would demonstrate its broader relevance to separation systems and reinforce its potential as a general framework for modelling, optimisation, and control of cyclic processes.

Dissemination of outputs

Journal Papers

4. **F. Michalopoulou**, E. Papadakis, and M. M. Papathanasiou. A reinforcement learning control strategy for chromatographic separation processes. *Submitted to Chemical Engineering Research and Design*, Under review. (Invited article in the *Emerging Stars II* special issue)

3. **F. Michalopoulou** and M. M. Papathanasiou. Integrated dual adaptive control of continuous chromatographic separation processes via reinforcement learning. *Journal of Chromatography A*, 1768, 466638, 2026. (Invited article in the *Advances in Preparative Chromatography* special issue)

2. **F. Michalopoulou** and M. M. Papathanasiou. An approach to hybrid modelling in chromatographic separation processes. *Digital Chemical Engineering*, 14:100215, 2025. (Invited article in the *Emerging Stars* special issue)

1. **F. Michalopoulou** and M. M. Papathanasiou. Assessment of data-driven modeling approaches for chromatographic separation processes. *AIChE Journal*, e18600, 2024. (Invited article in the *2024 Futures Issue*)

Peer-reviewed Conference Papers

1. **F. Michalopoulou** and M. M. Papathanasiou. Accelerated process optimization of chromatographic separations using a hybrid modeling approach. *IFAC-PapersOnLine*, 59(6):427–432, 2025.

Oral Conference Presentations

9. **F. Michalopoulou**, E. Papadakis, and M. M. Papathanasiou. Adaptive reinforcement learning control for chromatographic separation processes. *AIChE Annual Meeting*, 2025.

8. **F. Michalopoulou** and M. M. Papathanasiou. Monitoring and control of semi-continuous separation systems: A toolkit from white to black box approaches. *AIChE Annual Meeting*, 2025.

7. **F. Michalopoulou** and M. M. Papathanasiou. Reinforcement learning for automated control of semi-continuous chromatography. *ICB VII*, 2025.

6. **F. Michalopoulou** and M. M. Papathanasiou. Advanced modelling techniques for the optimisation of chromatographic separation processes. *ACS Spring*, 2025.

5. **F. Michalopoulou** and M. M. Papathanasiou. Towards grey-box modelling in chromatographic separation processes. *AIChE Annual Meeting*, 2024.

4. **F. Michalopoulou** and M. M. Papathanasiou. Machine learning aided tools for online monitoring of chromatographic separation processes. *AIChE Annual Meeting*, 2024.

3. **F. Michalopoulou** and M. M. Papathanasiou. From white- to black-box modelling in chromatographic separation processes. *SPICA*, 2024.

2. **F. Michalopoulou** and M. M. Papathanasiou. Data-driven modelling approaches for chromatographic separation processes. *ChemEngDayUK*, 2024.

1. **F. Michalopoulou** and M. M. Papathanasiou. From high-fidelity to data-driven modelling in separation systems: An application to chromatographic separations. *AIChE Annual Meeting*, 2023.

Bibliography

R. Agharafeie, J. R. C. Ramos, R. Oliveira, and J. M. Mendes. Application of Hybrid Models in Biopharmaceutical Manufacturing. *Chemie Ingenieur Technik*, 94(9):1281–1282, 2022. ISSN 0009-286X. doi: 10.1002/CITE.202255039.

R. Agharafeie, J. R. C. Ramos, J. M. Mendes, and R. Oliveira. From Shallow to Deep Bioprocess Hybrid Modeling: Advances and Future Perspectives. *Fermentation*, 9(10):922, 2023. ISSN 23115637. doi: 10.3390/FERMENTATION9100922/S1.

M. S. Alhajeri, Y. M. Ren, F. Ou, F. Abdullah, and P. D. Christofides. Model predictive control of nonlinear processes using transfer learning-based recurrent neural networks. *Chemical Engineering Research and Design*, 205:1–12, 2024. ISSN 0263-8762. doi: 10.1016/J.CHERD.2024.03.019.

K. Alhazmi, F. Albalawi, and S. M. Sarathy. A reinforcement learning-based economic model predictive control framework for autonomous operation of chemical reactors. *Chemical Engineering Journal*, 428:130993, 2022. ISSN 1385-8947. doi: 10.1016/J.CEJ.2021.130993.

R. D. Amos and R. Kobayashi. Feature Engineering for Materials Chemistry - Does Size Matter? *Journal of Chemical Information and Modeling*, 59(5):1873–1881, 2019. ISSN 1549960X. doi: 10.1021/ACS.JCIM.8B00977/ASSET/IMAGES/LARGE/CI-2018-009773{_}0013.JPEG.

D. Andersson, C. Edlund, B. Corbett, and R. Sjögren. Adaptable control policies for variable liquid chromatography columns using deep reinforcement learning. *Scientific Reports 2023 13:1*, 13(1):1–13, 2023. ISSN 2045-2322. doi: 10.1038/s41598-023-38145-2.

S. Andris and J. Hubbuch. Modeling of hydrophobic interaction chromatography for the separation of antibody-drug conjugates and its application towards quality by design. *Journal of biotechnology*, 317:48–58, 2020. ISSN 1873-4863. doi: 10.1016/J.JBIOTEC.2020.04.018.

A. Anupa, N. G. Jesubalan, R. Trivedi, N. Nitika, V. S. Buddhiraju, V. Runkana, and A. S. Rathore. Implementation of machine learning tool for continued process verification of process chromatography unit operation. *Journal of Chromatography A*, 1742:465642, 2025. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2024.465642.

A. Armstrong, K. Horry, T. Cui, M. Hulley, R. Turner, S. S. Farid, S. Goldrick, and D. G. Bracewell. Advanced control strategies for bioprocess chromatography: Challenges and opportunities for intensified processes and next generation products. *Journal of Chromatography A*, 1639:461914, 2021. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2021.461914.

L. Aumann and M. Morbidelli. A continuous multicolumn countercurrent solvent gradient purification (MCSGP) process. *Biotechnology and Bioengineering*, 98 (5):1043–1055, 2007. ISSN 00063592. doi: 10.1002/bit.21527.

F. Boukouvala, F. J. Muzzio, and M. G. Ierapetritou. Dynamic data-driven modeling of pharmaceutical processes. *Industrial and Engineering Chemistry Research*, 50(11):6743–6754, 2011. ISSN 08885885. doi: 10.1021/ie102305a.

K. W. Brhane, S. Qamar, and A. Seidel-Morgenstern. Two-Dimensional General Rate Model of Liquid Chromatography Incorporating Finite Rates of Adsorption-Desorption Kinetics and Core-Shell Particles. *Industrial and Engineering Chemistry Research*, 58(19):8296–8308, 2019. ISSN 15205045. doi: 10.1021/acs.iecr.9b00364.

E. Brochu, V. M. Cora, and N. de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. 2010. doi: 10.48550/arXiv.1012.2599.

L. Buşoniu, T. de Bruin, D. Tolić, J. Kober, and I. Palunko. Reinforcement learning for control: Performance, stability, and deep approximators. *Annual Reviews in Control*, 46:8–28, 2018. ISSN 1367-5788. doi: 10.1016/J.ARCONTROL.2018.09.005.

E. Cabo-Calvet, C. Ortiz-Bolsico, J. J. Baeza-Baeza, and M. C. García-Alvarez-Coque. Description of the Retention and Peak Profile for Chromolith Columns in Isocratic and Gradient Elution Using Mobile Phase Composition and Flow Rate as Factors. *Chromatography 2014, Vol. 1, Pages 194-210*, 1(4):194–210, 2014. ISSN 2227-9075. doi: 10.3390/CHROMATOGRAPHY1040194.

S. Chan, N. Titchener-Hooker, and E. Sørensen. Optimal economic design and operation of single- and multi-column chromatographic processes. *Biotechnology Progress*, 24(2):389–401, 2008. ISSN 87567938. doi: 10.1021/BP070270M,.

Y. Chen, O. Yang, C. Sampat, P. Bhalode, R. Ramachandran, and M. Ierapetritou. Digital Twins in Pharmaceutical and Biopharmaceutical Manufacturing: A Literature Review. *Processes 2020, Vol. 8, Page 1088*, 8(9):1088, 2020. ISSN 2227-9717. doi: 10.3390/PR8091088.

E. J. Close, J. R. Salm, D. G. Bracewell, and E. Sorensen. A model based approach for identifying robust operating conditions for industrial chromatography with process variability. *Chemical Engineering Science*, 116:284–295, 2014. ISSN 00092509. doi: 10.1016/J.CES.2014.03.010.

S. M. Cramer and M. A. Holstein. Downstream bioprocessing: recent advances and future promise. *Current Opinion in Chemical Engineering*, 1(1):27–37, 2011. ISSN 2211-3398. doi: 10.1016/J.COCHE.2011.08.008.

A. Creasy, G. Barker, Y. Yao, and G. Carta. Systematic interpolation method predicts protein chromatographic elution from batch isotherm data without a detailed mechanistic isotherm model. *Biotechnology Journal*, 10(9): 1400–1411, 2015. ISSN 18607314. doi: 10.1002/BIOT.201500089;PAGE: STRING:ARTICLE/CHAPTER.

J. Dai, P. W. Carr, and D. V. McCalley. A new approach to the determination of column overload characteristics in reversed-phase liquid chromatography. *Journal of Chromatography A*, 1216(12):2474–2482, 2009. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2009.01.036.

A. De Brébisson and P. Vincent. An Exploration of Softmax Alternatives Belonging to the Spherical Loss Family. *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, 2015. doi: doi.org/10.48550/arXiv.1511.05042.

M. Degerman, N. Jakobsson, and B. Nilsson. Constrained optimization of a preparative ion-exchange step for antibody purification. *Journal of Chromatography A*, 1113(1-2):92–100, 2006. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2006.01.121.

V. S. Devarakonda, W. Sun, X. Tang, and Y. Tian. Recent Advances in Reinforcement Learning for Chemical Process Control. *Processes 2025, Vol. 13, Page 1791*, 13(6):1791, 2025. ISSN 2227-9717. doi: 10.3390/PR13061791.

S. Dimartino, C. Boi, and G. C. Sarti. A validated model for the simulation of protein purification through affinity membrane chromatography. *Journal of Chromatography A*, 1218(13):1677–1690, 2011. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2010.11.056.

C. Ding and M. Ierapetritou. Machine learning-based optimization of a multi-step ion exchange chromatography for ternary protein separation. *Computers & Chemical Engineering*, 184:108642, 2024. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2024.108642.

C. Ding, C. Gerberich, and M. Ierapetritou. Hybrid model development for parameter estimation and process optimization of hydrophobic interaction chromatography. *Journal of Chromatography A*, 1703:464113, 2023. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2023.464113.

S. R. Dubey, S. K. Singh, and B. B. Chaudhuri. Activation functions in deep learning: A comprehensive survey and benchmark. *Neurocomputing*, 503: 92–108, 2022. ISSN 0925-2312. doi: 10.1016/J.NEUCOM.2022.06.111.

D. M. Ecker, S. D. Jones, and H. L. Levine. The therapeutic monoclonal antibody market, 2015. ISSN 19420870.

S. Engell and A. Toumi. Optimization and control of chromatography. *Computer Aided Chemical Engineering*, 18(C):9–22, 2004. ISSN 1570-7946. doi: 10.1016/S1570-7946(04)80081-6.

M. Enmark, J. Häggström, J. Samuelsson, and T. Fornstedt. Building machine-learning-based models for retention time and resolution predictions in ion pair chromatography of oligonucleotides. *Journal of Chromatography A*, 1671: 462999, 2022. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2022.462999.

T. Eslami and A. Jungbauer. Control strategy for biopharmaceutical production by model predictive control. *Biotechnology Progress*, 40(2):e3426, 2024. ISSN 1520-6033. doi: 10.1002/BTPR.3426.

A. Felinger and G. Guiochon. Comparison of the kinetic models of linear chromatography. *Chromatographia*, 60(SUPPL.):S175–S180, 2004. ISSN 00095893. doi: 10.1365/s10337-004-0288-7.

A. Felinger, A. Cavazzini, and F. Dondi. Equivalence of the microscopic and macroscopic models of chromatography: stochastic–dispersive versus lumped kinetic model. *Journal of Chromatography A*, 1043(2):149–157, 2004. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2004.05.081.

C. Feng, Q. Xu, X. Qiu, Y. Jin, J. Ji, Y. Lin, S. Le, J. She, D. Lu, and G. Wang. Evaluation and application of machine learning-based retention time prediction for suspect screening of pesticides and pesticide transformation products in LC-HRMS. *Chemosphere*, 271:129447, 2021. ISSN 0045-6535. doi: 10.1016/J.CHEMOSPHERE.2020.129447.

I. Fioretti, T. Müller-Späth, L. Aumann, and M. Sponchioni. UV-based dynamic control improves the robustness of multicolumn countercurrent solvent gradient purification of oligonucleotides. *Biotechnology Journal*, 19(7):2400170, 2024. ISSN 1860-7314. doi: 10.1002/BIOT.202400170.

A. Fuller, Z. Fan, C. Day, and C. Barlow. Digital Twin: Enabling Technologies, Challenges and Open Research. *IEEE Access*, 8:108952–108971, 2019. doi: 10.1109/ACCESS.2020.2998358.

J. R. Gardner, M. J. Kusner, Z. Xu, K. Q. Weinberger, and J. P. Cunningham. Bayesian Optimization with Inequality Constraints. In *31st International Conference on Machine Learning*, 2014.

X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. pages 249–256. JMLR Workshop and Conference Proceedings, 2010.

F. Gritti and G. Guiochon. Overload behavior and apparent efficiencies in chromatography. *Journal of Chromatography A*, 1254:30–42, 2012. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2012.07.015.

M. Grom, M. Kozorog, S. Caserman, A. Pohar, and B. Likozar. Protein A affinity chromatography of Chinese hamster ovary (CHO) cell culture broths containing biopharmaceutical monoclonal antibody (mAb): Experiments and mechanistic transport, binding and equilibrium modeling. *Journal of Chromatography B: Analytical Technologies in the Biomedical and Life Sciences*, 1083:44–56, 2018. ISSN 1873376X. doi: 10.1016/j.jchromb.2018.02.032.

C. Grossmann, G. Erdem, M. Morari, M. Amanullah, M. Mazzotti, and M. Morbidelli. 'Cycle to cycle' optimizing control of simulated moving beds. *AIChE Journal*, 54(1):194–208, 2008. ISSN 1547-5905. doi: 10.1002/AIC.11346.

C. Grossmann, G. Ströhlein, M. Morari, and M. Morbidelli. Optimizing model predictive control of the chromatographic multi-column solvent gradient purification (MCSGP) process. *Journal of Process Control*, 20(5):618–629, 2010. ISSN 0959-1524. doi: 10.1016/J.JPROCONT.2010.02.013.

G. Guiochon, A. Felinger, D. G. Shirazi, and A. M. Katti. *Fundamentals of Preparative and Nonlinearn Chromatography*. Academic Press, second edition, 2006. ISBN 9780123705372.

T. Hahn, T. Huuk, V. Heuveline, and J. Hubbuch. Simulating and Optimizing Preparative Protein Chromatography with ChromX. *Journal of Chemical Education*, 92(9):1497–1502, 2015. ISSN 19381328. doi: 10.1021/ed500854a.

T. Hahn, T. Huuk, A. Osberghaus, K. Doninger, S. Nath, S. Hepbildikler, V. Heuveline, and J. Hubbuch. Calibration-free inverse modeling of ion-exchange chromatography in industrial antibody purification. *Engineering in Life Sciences*, 16(2):107–113, 2016. ISSN 16182863. doi: 10.1002/elsc.201400248.

T. Hahn, N. Geng, K. Petrushevska-Seebach, M. E. Dolan, M. Scheindel, P. Graf, K. Takenaka, K. Izumida, L. Li, Z. Ma, and N. Schuelke. Mechanistic modeling, simulation, and optimization of mixed-mode chromatography for an antibody polishing step. *Biotechnology Progress*, 39(2):e3316, 2023. ISSN 15206033. doi: 10.1002/btpr.3316.

A. Hørsholt, L. H. Christiansen, K. Meyer, J. K. Huusom, and J. B. Jørgensen. A Discontinuous-Galerkin Finite-Element Method for Simulation of Packed Bed

Chromatographic Processes. *IFAC-PapersOnLine*, 52(1):346–351, 2019. ISSN 2405-8963. doi: 10.1016/J.IFACOL.2019.06.086.

D. R. Jones, M. Schonlau, and W. J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4): 455–492, 1998. ISSN 09255001. doi: 10.1023/A:1008306431147/METRICS.

V. S. Joshi, V. Kumar, and A. S. Rathore. Role of Organic Modifier and Gradient Shape in RP-HPLC Separation: Analysis of GCSF Variants. *Journal of Chromatographic Science*, 53(3):417–423, 2015. ISSN 0021-9665. doi: 10.1093/CHROMSCI/BMU222.

V. S. Joshi, V. Kumar, and A. S. Rathore. Optimization of ion exchange sigmoidal gradients using hybrid models: Implementation of quality by design in analytical method development. *Journal of Chromatography A*, 1491:145–152, 2017. ISSN 18733778. doi: 10.1016/j.chroma.2017.02.058.

A. Jungbauer, P. Satzer, A. Duerauer, A. Azevedo, R. Aires-Barros, B. Nilsson, S. Farid, S. Goldrick, M. Ottens, M. Sponchioni, and H. Marcelo Fernandez Lahore. Continuous downstream processing. *Separation and Purification Technology*, 338:126439, 2024. ISSN 1383-5866. doi: 10.1016/J.SEPPUR.2024.126439.

K. Kaczmarski and G. Guiochon. Modeling of the mass-transfer kinetics in chromatographic columns packed with shell and pellicular particles. *Analytical Chemistry*, 79(12):4648–4656, 2007. ISSN 00032700. doi: 10.1021/ac070209w.

K. Kaczmarski, D. Antos, H. Sajonz, P. Sajonz, and G. Guiochon. Comparative modeling of breakthrough curves of bovine serum albumin in anion-exchange chromatography. *Journal of Chromatography A*, 925(1-2):1–17, 2001. ISSN 0021-9673. doi: 10.1016/S0021-9673(01)01035-4.

S. J. Kalil, F. Maugeri, and M. I. Rodrigues. Response surface analysis and simulation as a tool for bioprocess design and optimization. *Process Biochemistry*, 35(6):539–550, 2000. ISSN 13595113. doi: 10.1016/S0032-9592(99)00101-6.

K. Katsoulas, M. Tirapelle, E. Sorensen, and L. Mazzei. On the apparent dispersion coefficient of the equilibrium dispersion model: An asymptotic analysis. *Journal of Chromatography A*, 1708:464345, 2023. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2023.464345.

K. Katsoulas, F. Galvanin, L. Mazzei, M. Besenhard, and E. Sorensen. Model-based design of experiments for efficient and accurate isotherm model identification in High Performance Liquid Chromatography. *Computers & Chemical Engineering*, 195:109021, 2025. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2025.109021.

Y. Kawajiri. Model-based optimization strategies for chromatographic processes: a review. 27:1–26, 2021. doi: 10.1007/s10450-020-00251-2.

Y. Kawajiri and L. T. Biegler. Large scale nonlinear optimization for asymmetric operation and design of Simulated Moving Beds. *Journal of Chromatography A*, 1133(1-2):226–240, 2006. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2006.08.037.

Y. Kawajiri and L. T. Biegler. Large scale optimization strategies for zone configuration of simulated moving beds. *Computers & Chemical Engineering*, 32(1-2):135–144, 2008. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2007.04.009.

S. Kay, H. Kay, A. W. Rogers, and D. Zhang. Integrating hybrid modelling and transfer learning for new bioprocess predictive modelling. *Computer Aided Chemical Engineering*, 52:2595–2600, 2023. ISSN 1570-7946. doi: 10.1016/B978-0-443-15274-0.50412-1.

A. Kensert, R. Bouwmeester, K. Efthymiadis, P. Van Broeck, G. Desmet, and D. Cabooter. Graph Convolutional Networks for Improved Prediction and Interpretability of Chromatographic Retention Data. *Analytical Chemistry*, 93(47):15633–15641, 2021. ISSN 15206882. doi: 10.1021/ACS.ANALCHEM.1C02988.

A. Kensert, E. Bosten, G. Collaerts, K. Efthymiadis, P. Van Broeck, G. Desmet, and D. Cabooter. Convolutional neural network for automated peak detection in reversed-phase liquid chromatography. *Journal of Chromatography A*, 1672:463005, 2022. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2022.463005.

A. Kensert, P. Libin, G. Desmet, and D. Cabooter. Deep reinforcement learning for the direct optimization of gradient separations in liquid chromatography. *Journal of Chromatography A*, 1720:464768, 2024. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2024.464768.

D. Keulen, E. van der Hagen, G. Geldhof, O. Le Bussy, M. Pabst, and M. Ottens. Using artificial neural networks to accelerate flowsheet optimization for downstream process development. *Biotechnology and Bioengineering*, 2023. ISSN 10970290. doi: 10.1002/BIT.28454.

O. Khanal and A. M. Lenhoff. Developments and opportunities in continuous biopharmaceutical manufacturing. *mAbs*, 13(1):1903664, 2021. ISSN 19420870. doi: 10.1080/19420862.2021.1903664.

E. J. Kim, J. H. Kim, M. S. Kim, S. H. Jeong, and D. H. Choi. Process Analytical Technology Tools for Monitoring Pharmaceutical Unit Operations: A Control Strategy for Continuous Process Verification. *Pharmaceutics 2021, Vol. 13, Page 919*, 13(6):919, 2021. ISSN 1999-4923. doi: 10.3390/PHARMACEUTICS13060919.

S. Kluters, F. Wittkopp, M. Jöhnck, and C. Frech. Application of linear pH gradients for the modeling of ion exchange chromatography: Separation of monoclonal antibody monomer from aggregates. *Journal of Separation Science*, 39(4):663–675, 2016. ISSN 16159314. doi: 10.1002/jssc.201500994.

O. Kohlbacher, S. Quinten, M. Sturm, B. M. Mayr, and C. G. Huber. Structure-activity relationships in chromatography: Retention prediction of oligonucleotides with support vector regression. *Angewandte Chemie - International Edition*, 45(42):7009–7012, 2006. ISSN 14337851. doi: 10.1002/ANIE.200602561.

M. Krättli, G. Ströhlein, L. Aumann, T. Müller-Späth, and M. Morbidelli. Closed loop control of the multi-column solvent gradient purification process. *Journal of Chromatography A*, 1218(50):9028–9036, 2011. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2011.09.081.

M. Krättli, T. Müller-Späth, and M. Morbidelli. Multifraction separation in countercurrent chromatography (MCSGP). *Biotechnology and Bioengineering*, 110(9):2436–2444, 2013a. ISSN 00063592. doi: 10.1002/bit.24901.

M. Krättli, F. Steinebach, and M. Morbidelli. Online control of the twin-column countercurrent solvent gradient process for biochromatography. *Journal of Chromatography A*, 1293:51–59, 2013b. ISSN 00219673. doi: 10.1016/j.chroma.2013.03.069.

V. Kumar and A. M. Lenhoff. Mechanistic Modeling of Preparative Column Chromatography for Biotherapeutics. *Annual Review of Chemical and Biomolecular Engineering*, 2020. doi: 10.1146/annurev-chembioeng.

V. Kumar, S. Leweke, W. Heymann, E. von Lieres, F. Schlegel, K. Westerberg, and A. M. Lenhoff. Robust mechanistic modeling of protein ion-exchange chromatography. *Journal of Chromatography A*, 1660:462669, 2021. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2021.462669.

A. Küpper and S. Engell. Engineering of Online Optimizing Control - A Case Study: Reactive SMB Chromatography. *IFAC Proceedings Volumes*, 41(2): 964–969, 2008. ISSN 1474-6670. doi: 10.3182/20080706-5-KR-1001.00165.

Y. A. LeCun, L. Bottou, G. B. Orr, and K. R. Müller. Neural Networks: Tricks of the Trade. In G. Montavon, G. B. Orr, and K. R. Müller, editors, *Lecture Notes in Computer Science*, volume 7700, chapter Efficient BackProp, pages 9–48. Springer, Berlin, Heidelberg, second edition, 2012. ISBN 978-3-642-35289-8. doi: 10.1007/978-3-642-35289-8{_}3.

J. W. Lee and A. Seidel-Morgenstern. Model Predictive Control of Simulated Moving Bed Chromatography for Binary and Pseudo-binary Separations: Simulation Study. *IFAC-PapersOnLine*, 51(18):530–535, 2018. ISSN 2405-8963. doi: 10.1016/J.IFACOL.2018.09.370.

S. Leweke and E. von Lieres. Chromatography Analysis and Design Toolkit (CADET). *Computers and Chemical Engineering*, 113:274–294, 2018. ISSN 00981354. doi: 10.1016/J.COMPCHEMENG.2018.02.025.

S. Li and S. Wang. A Method for Calculating Chromatographic Separation Process - Recurrent Neural Network. *SSRN*, 2024. doi: 10.2139/SSRN.5078440.

S. Li and S. Wang. Cascade recurrent neural network computation of chromatographic elution process. *Results in Chemistry*, 17:102583, 2025. ISSN 2211-7156. doi: 10.1016/J.RECHEM.2025.102583.

D. Q. Lin, Q. L. Zhang, and S. J. Yao. Model-assisted approaches for continuous chromatography: Current situation and challenges. *Journal of Chromatography A*, 1637, 2021. ISSN 18733778. doi: 10.1016/j.chroma.2020.461855.

S. Liu and L. G. Papageorgiou. Multi-objective optimisation for biopharmaceutical manufacturing under uncertainty. *Computers & Chemical Engineering*, pages 383–393, 2018. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2018.09.015.

S. Liu and L. G. Papageorgiou. Optimal Antibody Purification Strategies Using Data-Driven Models. *Engineering*, 5(6):1077–1092, 2019. ISSN 2095-8099. doi: 10.1016/J.ENG.2019.10.011.

C. D. Luca, S. Felletti, G. Lievore, A. Buratti, S. Vogg, M. Morbidelli, A. Cavazzini, M. Catani, M. Macis, A. Ricci, and W. Cabri. From batch to continuous chromatographic purification of a therapeutic peptide through multicolumn countercurrent solvent gradient purification. *Journal of Chromatography A*, 1625:461304, 2020. ISSN 18733778. doi: 10.1016/j.chroma.2020.461304.

M. K. Maruthamuthu, S. R. Rudge, A. M. Ardekani, M. R. Ladisch, and M. S. Verma. Process Analytical Technologies and Data Analytics for the Manufacture of Monoclonal Antibodies. *Trends in Biotechnology*, 38(10):1169–1186, 2020. ISSN 18793096. doi: 10.1016/J.TIBTECH.2020.07.004.

D. D. Matyushin, A. Y. Sholokhova, and A. K. Buryak. A deep convolutional neural network for the estimation of gas chromatographic retention indices. *Journal of Chromatography A*, 1607:460395, 2019. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2019.460395.

S. Mazraedoost, H. Sedigh Malekroodi, P. Žuvela, M. Yi, and J. J. Liu. Prediction of Chromatographic Retention Time of a Small Molecule from SMILES Representation Using a Hybrid Transformer-LSTM Model. *Journal of Chemical Information and Modeling*, 65(7):3343–3356, 2025. ISSN 1549960X. doi: 10.1021/ACS.JCIM.5C00167.

M. Mazzotti, G. Storti, and M. Morbidelli. Optimal operation of simulated moving bed units for nonlinear chromatographic separations. *Journal of Chromatography A*, 769(1):3–24, 1997. ISSN 0021-9673. doi: 10.1016/S0021-9673(97)00048-4.

K. McBride, E. I. Sanchez Medina, and K. Sundmacher. Hybrid Semi-parametric Modeling in Separation Processes: A Review. *Chemie Ingenieur Technik*, 92 (7):842–855, 2020. ISSN 1522-2640. doi: 10.1002/CITE.202000025.

G. Montúfar, R. Pascanu, K. Cho, and Y. Bengio. On the Number of Linear Regions of Deep Neural Networks. *Advances in Neural Information Processing Systems*, 4(January):2924–2932, 2014. ISSN 10495258. doi: doi.org/10.48550/arXiv.1402.1869.

M. Mouellef, F. L. Vetter, and J. Strube. Benefits and Limitations of Artificial Neural Networks in Process Chromatography Design and Operation. *Processes 2023, Vol. 11, Page 1115*, 11(4):1115, 2023. ISSN 2227-9717. doi: 10.3390/PR11041115.

T. Müller-Späth, L. Aumann, L. Melter, G. Ströhlein, and M. Morbidelli. Chromatographic separation of three monoclonal antibody variants using multicolumn countercurrent solvent gradient purification (MCSGP). *Biotechnology and Bioengineering*, 100(6):1166–1177, 2008. ISSN 00063592. doi: 10.1002/BIT.21843.

T. Müller-Späth, G. Ströhlein, L. Aumann, H. Kornmann, P. Valax, L. Delegrange, E. Charbaut, G. Baer, A. Lamproye, M. Jöhnck, M. Schulte, and M. Morbidelli. Model simulation and experimental verification of a cation-exchange IgG capture step in batch and continuous chromatography. *Journal of Chromatography A*, 1218(31):5195–5204, 2011. ISSN 00219673. doi: 10.1016/j.chroma.2011.05.103.

T. Müller-Späth, G. Ströhlein, O. Lyngberg, and D. Maclean. Enabling high purities and yields in therapeutic peptide purification using multicolumn countercurrent solvent gradient purification. *Chimica Oggi/Chemistry Today*, 31 (5):56–60, 2013.

D. Nagrath, A. Messac, B. W. Bequette, and S. M. Cramer. A Hybrid Model Framework for the Optimization of Preparative Chromatographic Processes. *Biotechnology Progress*, 20(1):162–178, 2004. ISSN 1520-6033. doi: 10.1021/BP034026G.

H. Narayanan, M. Luna, M. Sokolov, P. Arosio, A. Butté, and M. Morbidelli. Hybrid Models Based on Machine Learning and an Increasing Degree of Process Knowledge: Application to Capture Chromatographic Step. *Industrial and Engineering Chemistry Research*, 60(29):10466–10478, 2021a. ISSN 15205045. doi: 10.1021/ACS.IECR.1C01317.

H. Narayanan, T. Seidler, M. F. Luna, M. Sokolov, M. Morbidelli, and A. Butté. Hybrid Models for the simulation and prediction of chromatographic processes for protein capture. *Journal of Chromatography A*, 1650:462248, 2021b. ISSN 00219673. doi: 10.1016/j.chroma.2021.462248.

H. Narayanan, M. von Stosch, F. Feidl, M. Sokolov, M. Morbidelli, and A. Butté. Hybrid modeling for biopharmaceutical processes: advantages, opportunities, and implementation. *Frontiers in Chemical Engineering*, 5:1157889, 2023. ISSN 26732718. doi: 10.3389/FCENG.2023.1157889.

M. V. Narkhede, P. P. Bartakke, and M. S. Sutaone. A review on weight initialization strategies for neural networks. *Artificial Intelligence Review*, 55(1): 291–322, 2022. ISSN 15737462. doi: 10.1007/S10462-021-10033-Z/TABLES/15.

S. Natarajan and J. H. Lee. Repetitive model predictive control applied to a simulated moving bed chromatography system. *Computers & Chemical Engineering*, 24(2-7):1127–1133, 2000. ISSN 0098-1354. doi: 10.1016/S0098-1354(00)00493-2.

C. Ng, H. Osuna-Sanchez, E. Valéry, D. Bracewell, and E. Sørensen. Optimization of protein A chromatography for antibody capture. *Computer Aided Chemical Engineering*, 30:1367–1371, 2012. ISSN 1570-7946. doi: 10.1016/B978-0-444-59520-1.50132-9.

S. K. Niazi. Continuous Manufacturing of Recombinant Drugs: Comprehensive Analysis of Cost Reduction Strategies, Regulatory Pathways, and Global Implementation. *Pharmaceuticals*, 18(8):1157, 2025. ISSN 1424-8247. doi: 10.3390/PH18081157.

L. E. Niezen, P. J. Libin, D. Cabooter, and G. Desmet. Reinforcement learning for automated method development in liquid chromatography: insights in the reward scheme and experimental budget selection. *Journal of Chromatography A*, 1748:465845, 2025. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2025.465845.

S. Nikita, A. Tiwari, D. Sonawat, H. Kodamana, and A. S. Rathore. Reinforcement learning based optimization of process chromatography for continuous

processing of biopharmaceuticals. *Chemical Engineering Science*, 230, 2021. ISSN 00092509. doi: 10.1016/J.CES.2020.116171.

C. Oh, S. H. Zak, H. Mirzaei, C. Buck, F. E. Regnier, and X. Zhang. Neural network prediction of peptide separation in strong anion exchange chromatography. *Bioinformatics*, 23(1):114–118, 2007. ISSN 1367-4803. doi: 10.1093/BIOINFORMATICS/BTL561.

J. Orehek, D. Teslić, and B. Likozar. Continuous Crystallization Processes in Pharmaceutical Manufacturing: A Review. *Organic Process Research & Development*, 25(1):16–42, 2020. ISSN 1520586X. doi: 10.1021/ACS.OPRD.0C00398.

A. Osberghaus, S. Hepbildikler, S. Nath, M. Haindl, E. von Lieres, and J. Hubbuch. Optimizing a chromatographic three component separation: A comparison of mechanistic and empiric modeling approaches. *Journal of Chromatography A*, 1237:86–95, 2012. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2012.03.029.

C. C. Pantelides and F. E. Pereira. The future of digital applications in pharmaceutical operations. *Current Opinion in Chemical Engineering*, 45:101038, 2024. ISSN 2211-3398. doi: 10.1016/J.COCHE.2024.101038.

M. M. Papathanasiou, S. Avraamidou, R. Oberdieck, A. Mantalaris, F. Steinebach, M. Morbidelli, T. Mueller-Spaeth, and E. N. Pistikopoulos. Advanced control strategies for the multicolumn countercurrent solvent gradient purification process. *AIChE Journal*, 62(7):2341–2357, 2016. ISSN 15475905. doi: 10.1002/aic.15203.

M. M. Papathanasiou, B. Burnak, J. Katz, N. Shah, and E. N. Pistikopoulos. Assisting continuous biomanufacturing through advanced control in downstream purification. *Computers and Chemical Engineering*, 125:232–248, 2019a. ISSN 00981354. doi: 10.1016/j.compchemeng.2019.03.013.

M. M. Papathanasiou, B. Burnak, J. Katz, N. Shah, and E. N. Pistikopoulos. Control of a dual mode separation process via multi-parametric Model Predictive Control. *IFAC-PapersOnLine*, 52(1):988–993, 2019b. ISSN 2405-8963. doi: 10.1016/J.IFACOL.2019.06.191.

P. Petsagkourakis, I. O. Sandoval, E. Bradford, D. Zhang, and E. A. del Rio-Chanona. Reinforcement learning for batch bioprocess optimization. *Computers & Chemical Engineering*, 133:106649, 2020. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2019.106649.

D. Pfister, L. David, M. Holzer, and R. M. Nicoud. Designing affinity chromatographic processes for the capture of antibodies. Part I: A simplified approach. *Journal of Chromatography A*, 1494:27–39, 2017. ISSN 18733778. doi: 10.1016/j.chroma.2017.02.070.

T. D. Pham, C. Manapragada, Y. Sun, R. Bassett, and U. Aickelin. A scoping review of supervised learning modelling and data-driven optimisation in monoclonal antibody process development. *Digital Chemical Engineering*, 7: 100080, 2023. ISSN 2772-5081. doi: 10.1016/J.DCHE.2022.100080.

W. Pia tkowski, D. Antos, and K. Kaczmarski. Modeling of preparative chromatography processes with slow intraparticle mass transport kinetics. *Journal of Chromatography A*, 988(2):219–231, 2003. ISSN 0021-9673. doi: 10.1016/S0021-9673(02)02060-5.

S. M. Pirrung, L. A. van der Wielen, R. F. van Beckhoven, E. J. van de Sandt, M. H. Eppink, and M. Ottens. Optimization of biopharmaceutical downstream processes supported by mechanistic models and artificial neural networks. *Biotechnology Progress*, 33(3):696–707, 2017. ISSN 1520-6033. doi: 10.1002/BTPR.2435.

Y. Qu, I. Baker, J. Black, L. Fabri, S. L. Gras, A. M. Lenhoff, and S. E. Kentish. Application of mechanistic modelling in membrane and fiber chromatography for purification of biotherapeutics — A review. *Journal of Chromatography A*, 1716:464588, 2024. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2023.464588.

A. Rajendran, G. Paredes, and M. Mazzotti. Simulated moving bed chromatography for the separation of enantiomers. *Journal of Chromatography A*, 1216 (4):709–738, 2009. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2008.10.075.

C. E. Rasmussen and C. K. I. Williams. Gaussian Processes for Machine Learning. *Gaussian Processes for Machine Learning*, 2005. doi: 10.7551/MITPRESS/3206.001.0001.

A. S. Rathore, S. Nikita, G. Thakur, and S. Mishra. Artificial intelligence and machine learning applications in biopharmaceutical manufacturing. *Trends in Biotechnology*, 41(4):497–510, 2023. ISSN 0167-7799. doi: 10.1016/J.TIBTECH.2022.08.007.

M. J. Regufe, V. V. Santana, A. F. Ferreira, A. M. Ribeiro, J. M. Loureiro, and I. B. Nogueira. A Hybrid Modeling Framework for Membrane Separation Processes: Application to Lithium-Ion Recovery from Batteries. *Processes 2021, Vol. 9, Page 1939*, 9(11):1939, 2021. ISSN 2227-9717. doi: 10.3390/PR9111939.

J. U. Rehman, A. Muneer, and S. Qamar. Analysis of Equilibrium Dispersive Model of Liquid Chromatography Considering a Quadratic-Type Adsorption Isotherm. *Thermal Science*, 26(3A):2069–2080, 2022. doi: 10.2298/TSCI201229179U.

N. Rehman and S. Qamar. Numerical Approximation of the Nonequilibrium Model of Gradient Elution Chromatography Considering Linear and Nonlinear Solvent Strength Models. *ACS Omega*, 7(36):31905–31915, 2022. ISSN 24701343. doi: 10.1021/acsomega.2c02754.

M. Rüdt, N. Brestrich, L. Rolinger, and J. Hubbuch. Real-time monitoring and control of the load phase of a protein A capture step. *Biotechnology and Bioengineering*, 114(2):368–373, 2017. ISSN 10970290. doi: 10.1002/bit.26078.

J. Sansana, R. Rendall, I. Castillo, L. Chiang, and M. S. Reis. Hybrid modeling for improved extrapolation and transfer learning in the chemical processing industry. *Chemical Engineering Science*, 300:120568, 2024. ISSN 0009-2509. doi: 10.1016/J.CES.2024.120568.

V. V. Santana, M. S. Gama, J. M. Loureiro, A. E. Rodrigues, A. M. Ribeiro, F. W. Tavares, A. G. Barreto, and I. B. Nogueira. A First Approach towards Adsorption-Oriented Physics-Informed Neural Networks: Monoclonal Antibody Adsorption Performance on an Ion-Exchange Column as a Case Study. *ChemEngineering 2022, Vol. 6, Page 21*, 6(2):21, 2022. ISSN 2305-7084. doi: 10.3390/CHEMENGINEERING6020021.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. K. Openai. Proximal Policy Optimization Algorithms. *ArXiv*, 2017. doi: 10.48550/arXiv.1707.06347.

B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. ISSN 15582256. doi: 10.1109/JPROC.2015.2494218.

D. A. Shastri, K. Dhumal, A. Patil, D. Sawant, and D. S. Hatolkar. Process Analytical Technology (PAT): Enhancing Quality and Efficiency in Pharmaceutical Development and Production. *Journal for ReAttach Therapy and Developmental Diversities*, 6(1):1963–1973, 2023. ISSN 2589-7799. doi: 10.53555/JRTDD.V6I1.3112.

L. K. Shekhawat and A. S. Rathore. An overview of mechanistic modeling of liquid chromatography. *Preparative Biochemistry and Biotechnology*, 49(6):623–638, 2019. ISSN 15322297. doi: 10.1080/10826068.2019.1615504.

T. C. Silva, M. Eppink, and M. Ottens. Digital twin in high throughput chromatographic process development for monoclonal antibodies. *Journal of Chromatography A*, 1717:464672, 2024. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2024.464672.

J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian Optimization of Machine Learning Algorithms. *Advances in Neural Information Processing Systems*, 4:2951–2959, 2012. ISSN 10495258. doi: 10.48550/arXiv.1206.2944.

I. M. Sobol' and B. V. Shukman. Random and quasirandom sequences: Numerical estimates of uniformity of distribution. *Mathematical and Computer Modelling*, 18(8):39–45, 1993. ISSN 0895-7177. doi: 10.1016/0895-7177(93)90160-Z.

D. Solle, B. Hitzmann, C. Herwig, M. Pereira Remelhe, S. Ulonska, L. Wuerth, A. Prata, and T. Steckenreiter. Between the Poles of Data-Driven and Mechanistic Modeling for Process Operation. *Chemie Ingenieur Technik*, 89(5):542–561, 2017. ISSN 1522-2640. doi: 10.1002/CITE.201600175.

D. Solomatine, L. See, and R. Abrahart. Data-Driven Modelling: Concepts, Approaches and Experiences. In *Practical Hydroinformatics*, pages 17–30. Springer Berlin Heidelberg, 2008. doi: 10.1007/978-3-540-79881-1{_}2.

S. Sommerfeld and J. Strube. Challenges in biotechnology production - Generic processes and process optimization for monoclonal antibodies. *Chemical Engineering and Processing: Process Intensification*, 44(10):1123–1137, 2005. ISSN 02552701. doi: 10.1016/j.cep.2005.03.006.

F. Steinebach, M. Angarita, D. J. Karst, T. Müller-Späth, and M. Morbidelli. Model based adaptive control of a continuous capture process for monoclonal antibodies production. *Journal of Chromatography A*, 1444:50–56, 2016. ISSN 18733778. doi: 10.1016/J.CHROMA.2016.03.014.

F. Steinebach, N. Ulmer, L. Decker, L. Aumann, and M. Morbidelli. Experimental design of a twin-column countercurrent gradient purification process. *Journal of Chromatography A*, 1492:19–26, 2017. ISSN 18733778. doi: 10.1016/j.chroma.2017.02.049.

J. Strube, F. Grote, and R. Ditz. Bioprocess Design and Production Technology for the Future. In *Biopharmaceutical Production Technology, Volume 1 & Volume 2*, volume 2, pages 657–705. Wiley-VCH, 2012. ISBN 9783527330294. doi: 10.1002/9783527653096.ch20.

S. G. Subraveti, Z. Li, V. Prasad, and A. Rajendran. Can a computer "learn" nonlinear chromatography?: Physics-based deep neural networks for simulation and optimization of chromatographic processes. *Journal of Chromatography A*, 1672:463037, 2022a. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2022.463037.

S. G. Subraveti, Z. Li, V. Prasad, and A. Rajendran. Physics-Based Neural Networks for Simulation and Synthesis of Cyclic Adsorption Processes. *Industrial and Engineering Chemistry Research*, 61(11):4095–4113, 2022b. ISSN 15205045. doi: 10.1021/ACS.IECR.1C04731.

Y. N. Sun, W. W. Chen, S. J. Yao, and D. Q. Lin. Model-assisted process development, characterization and design of continuous chromatography for antibody separation. *Journal of Chromatography A*, 1707:464302, 2023. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2023.464302.

P. Suvarov, A. V. Wouwer, and A. Kienle. A simple robust control for simulated moving bed chromatographic separation. *IFAC Proceedings Volumes*, 45(15): 137–142, 2012. ISSN 1474-6670. doi: 10.3182/20120710-4-SG-2026.00098.

S. Taguado Menza, R. Prestia, I. Fioretti, and M. Sponchioni. Model-based optimization strategy for intensification in the chromatographic purification of oligonucleotides. *Journal of Chromatography A*, 1736:465321, 2024. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2024.465321.

S. Y. Tang, Y. H. Yuan, Y. C. Chen, S. J. Yao, Y. Wang, and D. Q. Lin. Physics-informed neural networks to solve lumped kinetic model for chromatography process. *Journal of Chromatography A*, 1708:464346, 2023. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2023.464346.

H. K. Teoh, M. Turner, N. Titchener-Hooker, and E. Sorensen. Experimental verification and optimisation of a detailed dynamic high performance liquid chromatography column model. *Computers & Chemical Engineering*, 25(4-6): 893–903, 2001. ISSN 0098-1354. doi: 10.1016/S0098-1354(01)00654-8.

A. Tiwari, V. S. Masampally, A. Agarwal, and A. S. Rathore. Digital twin of a continuous chromatography process for mAb purification: Design and model-based control. *Biotechnology and Bioengineering*, 120(3):748–766, 2023. ISSN 10970290. doi: 10.1002/BIT.28307,.

C. Ündey, S. Ertun, T. Mistretta, and B. Looze. Applied advanced process analytics in biopharmaceutical manufacturing: Challenges and prospects in real-time monitoring and control. *Journal of Process Control*, 20(9):1009–1018, 2010. ISSN 0959-1524. doi: 10.1016/J.JPROCONT.2010.05.008.

U.S. Food and Drug Administration (FDA). Guidance for Industry: PAT — A Framework for Innovative Pharmaceutical Development, Manufacturing, and Quality Assurance. Technical report, U.S. Department of Health and Human Services, Food and Drug Administration, Center for Drug Evaluation and Research (CDER), Rockville, MD, 2004.

M. Vaškevičius, J. Kapočiūtė-Dzikienė, and L. Šlepikas. Prediction of Chromatography Conditions for Purification in Organic Synthesis Using Deep Learning. *Molecules*, 26(9):2474, 2021. ISSN 14203049. doi: 10.3390/MOLECULES26092474.

L. Vetter-Joss, A. Sellberg, and E. B. Hansen. Model-Supported Design of Multicolumn Countercurrent Solvent Gradient Purification Processes for Rapid Implementation. *Industrial and Engineering Chemistry Research*, 63(41): 17612–17624, 2024. ISSN 15205045. doi: 10.1021/acs.iecr.4c02092.

E. von Lieres and J. Andersson. A fast and accurate solver for the general rate model of column liquid chromatography. *Computers & Chemical Engineering*, 34(8):1180–1191, 2010. ISSN 0098-1354. doi: 10.1016/J.COMPCHEMENG.2010.03.008.

G. Walsh. Biopharmaceutical benchmarks 2018. *Nature Biotechnology*, 36(12): 1136–1145, 2018. ISSN 15461696. doi: 10.1038/NBT.4305;KWRD.

C. Wang, K. U. Klatt, G. Dünnebier, S. Engell, and F. Hanisch. Neural network-based identification of SMB chromatographic processes. *Control Engineering Practice*, 11(8):949–959, 2003. ISSN 0967-0661. doi: 10.1016/S0967-0661(02)00212-5.

G. Wang, T. Briskot, T. Hahn, P. Baumann, and J. Hubbuch. Estimation of adsorption isotherm and mass transfer parameters in protein chromatography using artificial neural networks. *Journal of Chromatography A*, 1487:211–217, 2017a. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2017.01.068.

G. Wang, T. Briskot, T. Hahn, P. Baumann, and J. Hubbuch. Root cause investigation of deviations in protein chromatography based on mechanistic models and artificial neural networks. *Journal of Chromatography A*, 1515: 146–153, 2017b. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2017.07.089.

A. Ward and R. Pini. Efficient Bayesian Optimization of Industrial-Scale Pressure-Vacuum Swing Adsorption Processes for CO2 Capture. *Industrial and Engineering Chemistry Research*, 61(36):13650–13668, 2022. ISSN 15205045. doi: 10.1021/ACS.IECR.2C02313.

R. Webb, P. Doble, and M. Dawson. Optimisation of HPLC gradient separations using artificial neural networks (ANNs): Application to benzodiazepines in post-mortem samples. *Journal of Chromatography B*, 877(7):615–620, 2009. ISSN 1570-0232. doi: 10.1016/J.JCHROMB.2009.01.012.

K. Wong, R. Dornberger, and T. Hanne. An analysis of weight initialization methods in connection with different activation functions forfeedforward neural networks. *Evolutionary Intelligence*, 17(3):2081–2089, 2024. ISSN 18645917. doi: 10.1007/S12065-022-00795-Y/FIGURES/9.

C. F. Xie, H. Zhang, and R. C. Hwang. Application of Intelligent Control in Chromatography Separation Process. *Processes 2023, Vol. 11, Page 3443*, 11 (12):3443, 2023. ISSN 2227-9717. doi: 10.3390/PR11123443.

C. T. Yang, E. Kristiani, Y. K. Leong, and J. S. Chang. Big data and machine learning driven bioprocessing – Recent trends and critical analysis. *Bioresource Technology*, 372:128625, 2023. ISSN 0960-8524. doi: 10.1016/J.BIORTECH.2023.128625.

Y. X. Yang, Y. C. Chen, S. J. Yao, and D. Q. Lin. Parameter-by-parameter estimation method for adsorption isotherm in hydrophobic interaction chromatography. *Journal of Chromatography A*, 1716:464638, 2024. ISSN 0021-9673. doi: 10.1016/J.CHROMA.2024.464638.

Y. Yifei and S. Lakshminarayanan. Multi-agent reinforcement learning for process control: Exploring the intersection between fields of reinforcement learning, control theory, and game theory. *The Canadian Journal of Chemical Engineering*, 101(11):6227–6239, 2023. ISSN 1939-019X. doi: 10.1002/CJCE.24878.

X. Zhao, D. Zhang, L. Han, T. Zhang, and B. Xu. ODE-based Recurrent Model-free Reinforcement Learning for POMDPs. *NeurIPS Proceedings*, 2023. doi: 10.48550/arXiv.2309.14078.

Appendix A

Neural network architectures

A.1 Multiple-input/multiple-output networks

Table A.1 presents the optimal hyperparameters of the multiple-input/multiple-output (MIMO) artificial neural networks (ANNs) developed in Chapter 4 for predicting the elution concentrations of all species, as identified through Bayesian optimisation (BO).

Table A.1: Multiple-input/multiple-output artificial neural network architectures for predicting the elution concentrations of all species.

Activation function	**Training sample size**	**Hidden layers**	**Neurons per hidden layer**	**Learning rate**
ReLU	50 Cycles	3	83	$5.44 \cdot 10^{-3}$
	100 Cycles	2	124	$2.73 \cdot 10^{-3}$
	200 Cycles	3	195	$2.22 \cdot 10^{-3}$
	400 Cycles	3	161	$4.26 \cdot 10^{-3}$
Sigmoid	50 Cycles	3	82	$2.87 \cdot 10^{-2}$
	100 Cycles	3	37	$2.74 \cdot 10^{-2}$
	200 Cycles	4	70	$3.52 \cdot 10^{-2}$
	400 Cycles	2	199	$2.61 \cdot 10^{-2}$
Softmax	50 Cycles	3	47	$3.13 \cdot 10^{-2}$
	100 Cycles	2	105	$5.14 \cdot 10^{-2}$
	200 Cycles	2	96	$8.63 \cdot 10^{-2}$
	400 Cycles	2	150	$9.14 \cdot 10^{-3}$
Tanh	50 Cycles	3	107	$2.39 \cdot 10^{-3}$
	100 Cycles	1	32	$8.52 \cdot 10^{-3}$
	200 Cycles	3	45	$5.17 \cdot 10^{-3}$
	400 Cycles	2	251	$2.31 \cdot 10^{-3}$

A.2 Multiple-input/single-output networks

Table A.2 presents the optimal hyperparameters of the multiple-input/single-output (MISO) ANNs developed in Chapter 4 for predicting the elution concentration of the modifier, as identified through BO.

Table A.2: Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the modifier.

Activation function	**Training sample size**	**Hidden layers**	**Neurons per hidden layer**	**Learning rate**
ReLU	50 Cycles	2	126	$2.17 \cdot 10^{-3}$
	100 Cycles	2	76	$2.19 \cdot 10^{-3}$
	200 Cycles	2	185	$3.27 \cdot 10^{-3}$
	400 Cycles	2	145	$2.53 \cdot 10^{-3}$
Sigmoid	50 Cycles	3	133	$5.41 \cdot 10^{-3}$
	100 Cycles	2	252	$4.99 \cdot 10^{-3}$
	200 Cycles	3	91	$8.71 \cdot 10^{-3}$
	400 Cycles	3	135	$6.79 \cdot 10^{-3}$
Softmax	50 Cycles	2	162	$4.88 \cdot 10^{-3}$
	100 Cycles	3	54	$9.66 \cdot 10^{-3}$
	200 Cycles	2	65	$6.25 \cdot 10^{-3}$
	400 Cycles	2	34	$7.00 \cdot 10^{-3}$
Tanh	50 Cycles	2	54	$1.05 \cdot 10^{-3}$
	100 Cycles	2	234	$4.16 \cdot 10^{-4}$
	200 Cycles	2	143	$5.48 \cdot 10^{-3}$
	400 Cycles	2	177	$3.23 \cdot 10^{-3}$

Table A.3 presents the optimal hyperparameters of the MISO ANNs developed in Chapter 4 for predicting the elution concentration of the weak impurities, as identified through BO.

Table A.3: Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the weak impurities.

Activation function	**Training sample size**	**Hidden layers**	**Neurons per hidden layer**	**Learning rate**
ReLU	50 Cycles	1	172	$9.04 \cdot 10^{-3}$
	100 Cycles	2	68	$3.09 \cdot 10^{-3}$
	200 Cycles	1	139	$3.47 \cdot 10^{-3}$
	400 Cycles	3	172	$6.01 \cdot 10^{-3}$
Sigmoid	50 Cycles	3	89	$7.87 \cdot 10^{-3}$
	100 Cycles	4	55	$1.00 \cdot 10^{-2}$
	200 Cycles	4	60	$1.00 \cdot 10^{-2}$
	400 Cycles	3	159	$7.65 \cdot 10^{-3}$
Softmax	50 Cycles	2	198	$8.66 \cdot 10^{-3}$
	100 Cycles	2	221	$5.77 \cdot 10^{-3}$
	200 Cycles	2	82	$5.20 \cdot 10^{-3}$
	400 Cycles	1	68	$8.88 \cdot 10^{-3}$
Tanh	50 Cycles	1	134	$2.36 \cdot 10^{-3}$
	100 Cycles	3	32	$5.41 \cdot 10^{-3}$
	200 Cycles	3	92	$5.75 \cdot 10^{-3}$
	400 Cycles	3	85	$2.17 \cdot 10^{-3}$

Table A.4 presents the optimal hyperparameters of the MISO ANNs developed in Chapter 4 for predicting the elution concentration of the product, as identified through BO.

Table A.4: Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the product.

Activation function	**Training sample size**	**Hidden layers**	**Neurons per hidden layer**	**Learning rate**
ReLU	50 Cycles	3	81	$3.10 \cdot 10^{-3}$
	100 Cycles	3	194	$4.07 \cdot 10^{-3}$
	200 Cycles	2	170	$4.47 \cdot 10^{-3}$
	400 Cycles	3	241	$5.36 \cdot 10^{-3}$
Sigmoid	50 Cycles	3	44	$2.77 \cdot 10^{-3}$
	100 Cycles	4	50	$1.69 \cdot 10^{-2}$
	200 Cycles	3	76	$9.95 \cdot 10^{-3}$
	400 Cycles	3	167	$7.26 \cdot 10^{-3}$
Softmax	50 Cycles	1	160	$1.00 \cdot 10^{-2}$
	100 Cycles	2	106	$6.21 \cdot 10^{-3}$
	200 Cycles	2	232	$8.36 \cdot 10^{-3}$
	400 Cycles	1	173	$7.30 \cdot 10^{-3}$
Tanh	50 Cycles	2	40	$3.23 \cdot 10^{-3}$
	100 Cycles	3	110	$1.19 \cdot 10^{-3}$
	200 Cycles	2	104	$2.86 \cdot 10^{-3}$
	400 Cycles	3	76	$1.56 \cdot 10^{-3}$

Table A.5 presents the optimal hyperparameters of the MISO ANNs developed in Chapter 4 for predicting the elution concentration of the strong impurities, as identified through BO.

Table A.5: Multiple-input/single-output artificial neural network architectures for predicting the elution concentration of the strong impurities.

Activation function	**Training sample size**	**Hidden layers**	**Neurons per hidden layer**	**Learning rate**
ReLU	50 Cycles	2	50	$2.35 \cdot 10^{-3}$
	100 Cycles	1	171	$6.61 \cdot 10^{-3}$
	200 Cycles	1	155	$2.69 \cdot 10^{-3}$
	400 Cycles	4	211	$1.63 \cdot 10^{-3}$
Sigmoid	50 Cycles	4	70	$1.00 \cdot 10^{-2}$
	100 Cycles	4	51	$1.00 \cdot 10^{-2}$
	200 Cycles	3	84	$5.83 \cdot 10^{-3}$
	400 Cycles	3	148	$7.13 \cdot 10^{-3}$
Softmax	50 Cycles	1	187	$8.76 \cdot 10^{-2}$
	100 Cycles	1	128	$3.29 \cdot 10^{-3}$
	200 Cycles	2	182	$5.98 \cdot 10^{-3}$
	400 Cycles	2	85	$4.79 \cdot 10^{-3}$
Tanh	50 Cycles	3	120	$1.51 \cdot 10^{-3}$
	100 Cycles	4	32	$5.59 \cdot 10^{-3}$
	200 Cycles	4	52	$1.33 \cdot 10^{-3}$
	400 Cycles	2	166	$1.30 \cdot 10^{-3}$