

Western Australian School of Mines

**Monitoring and Control of Hydrocyclones by Use of Convolutional
Neural Networks and Deep Reinforcement Learning**

Keith Carmelo Giglia

0000-0002-5606-4199

This thesis is presented for the Degree of

Doctor of Philosophy

of

Curtin University

October 2022

Declaration

To the best of my knowledge and belief, this thesis contains no material previously published by any other person except where due acknowledgment has been made.

This thesis contains no material which has been accepted for the award of any other degree or diploma in any university.

Keith Giglia

9/10/22

Date

Abstract

The controlled operation of hydrocyclones in mineral processing circuits leads to improved efficiency and cost benefits. In industry, hydrocyclones are typically only monitored at the cluster level and use low-level control for stable operation. Sensing options to monitor individual hydrocyclone operation and product quality have been developed, along with advanced control methods. However, they have not found wide-spread industrial implementation. This research project explored image-based hydrocyclone performance monitoring using convolutional neural networks and hydrocyclone underflow video footage. Coupled with this advanced level of process monitoring, model-free reinforcement learning was explored for process control performance benefits.

Proof-of-concept investigations were performed using pretrained convolutional neural networks to produce hydrocyclone operational state classifiers on footage from a fixed high-speed camera filming a laboratory hydrocyclone and varied footage from multiple sources. A degree of robustness to the types of visual challenges that would be experienced in an industrial system were indicated. Further hydrocyclone operational state detector development was performed using fixed camera industrial footage, complicated by the lack of applied lighting. The trained classifier model's individual frame performance deteriorated on transition operating state periods and crop positions outside that used in training. It did however present strong discernment between correct fan state operation and the fault states (roping or blocked) under industrial conditions. Applied sensor considerations and the development of a health check to monitor for obstruction or camera movement was also presented.

Hydrocyclone underflow particle size inference from video footage was also investigated. Initial test work was performed using high-speed footage of a laboratory hydrocyclone and fine-tuning a pretrained convolutional neural network for the new task of outputting a particle size estimate. The discernment between fine and coarse particle sizes of the test dataset was sufficient to warrant further investigation. Standard frame rate video footage and operating data from an industrial hydrocyclone circuit was then collected. Neural network based models for

hydrocyclone underflow particle size prediction were developed from images only, circuit sensor data only, along with combining images and sensor data. The image-based models using convolutional neural networks showed stronger test dataset performance than the sensor only models. The combination of image and sensor data in the model architecture did not yield a further performance improvement under the scenario trialled.

A reinforcement learning algorithm, without access to a process model, was trialled interacting with a simulated open circuit hydrocyclone system. The operating data available to the reinforcement learning controller was that typically measured in such circuits, along with hydrocyclone operating state and underflow particle size information. Three continuous actions were available for hydrocyclone control, and with a meaningful reward function devised along with safety limits considered, the controller's behaviour was explored under several scenarios. The inclusion of historic sensor readings was used to compensate for incomplete state information and sensor noise. The performance could deteriorate under certain scenarios, highlighting that control override protections would be prudent to include in applied industrial control. The controller was able to learn to avoid the hydrocyclone roping fault state and operate near limits when the set points were not achievable. The controller was able to generalise to conditions outside those experienced in training and handle set point changes.

Through this work, proof-of-concept and applied considerations for convolutional neural network based hydrocyclone state detector and underflow particle size inference sensors were demonstrated. This forms the foundation for the development of industrial online monitoring of hydrocyclones through image analysis using this method. The behaviour of model-free reinforcement learning in a mineral processing context, incorporating additional information provided by the sensors developed, was found to be desirable. Building upon the practical considerations addressed, the future production of a safe and stable industrial reinforcement learning method would be beneficial for advanced circuit control.

Acknowledgements

Firstly, I would like to thank my supervisor Professor Chris Aldrich for his support throughout this journey. BHP and Brent Sullivan, Minara Resources and Michael Down, along with the anonymous industrial partner for the donation of hydrocyclone footage. IGO, Matthew Spagnolo, Paul Hudson, and the Nova mine site team for the donation of hydrocyclone footage and assistance during onsite data collection. Dr Boris Albijanic, Sheree McCutcheon, and the WASM Kalgoorlie lab staff for their assistance during laboratory hydrocyclone experiments. Alex Ciluzzo and Robert Kochmanski for letting me bounce ideas off them. Ben Lau and Molycop for allowing me to build upon their software. The Minerals Research Institute of Western Australia and Curtin University for their financial and other support for this project. NVIDIA Corporation for the donation of a Titan Xp GPU used in this research. A special thanks to the Curtin University library and copyright staff whose assistance has been invaluable. And finally, to my wife Patrícia and son Tiago for their patience during this time.

Copyright

I warrant that I have obtained, where necessary, permission from the copyright owners to use any third-party copyright material reproduced in the thesis, or to use any of my own published work in which the copyright is held by another party.

Publication

This thesis contains material from the below co-authored publication; predominantly in Section 2.3.4, Section 3.1, Section 3.2, and Appendix A along with other use throughout as cited.

<https://doi.org/10.1016/j.mineng.2020.106211>

This article was published in Minerals Engineering, Vol 149, K. C. Giglia & C. Aldrich, Operational state detection in hydrocyclones with convolutional neural networks and transfer learning, Article 106211, pp. 1–9, Copyright Elsevier (2020).

Attribution Statement

	Conception and Design	Acquisition of Data & Method	Data Conditioning & Manipulation	Analysis & Statistical Method	Interpretation & Discussion	Final Approval
Keith Giglia	X	X	X	X	X	X
I acknowledge that these represent my contribution to the above research output Signed.						
Prof Chris Aldrich	X					X
I acknowledge that these represent my contribution to the above research output Signed.						

Table of Contents

Declaration	i
Abstract	ii
Acknowledgements.....	iv
Copyright.....	v
Publication.....	v
Table of Contents	vi
List of Abbreviations and Acronyms	x
List of Symbols.....	xii
List of Figures	xv
List of Tables.....	xx
1 Introduction	1
1.1 Background	1
1.2 Objectives.....	2
1.3 Thesis Outline.....	3
2 Literature Review	4
2.1 Hydrocyclone Basics.....	4
2.2 Empirical Models.....	8
2.3 Operational State Detection	10
2.3.1 Electrical Tomographic Methods	11
2.3.2 Acoustic Methods	12
2.3.3 Underflow Probe Methods	13
2.3.4 Optical Methods.....	14
2.3.5 Summary	16
2.4 Online Particle Size Monitoring	16
2.4.1 Sub-sample Measurement.....	17
2.4.2 Soft Sensors.....	17
2.4.3 Overflow Discharge	18
2.4.4 Underflow Discharge.....	18
2.4.5 Summary	20

2.5	Convolutional Neural Networks.....	21
2.5.1	Background	21
2.5.2	Notable Networks	27
2.5.3	Mineral Processing.....	29
2.5.4	Summary	31
2.6	Hydrocyclone Control	31
2.6.1	Low-Level Control	31
2.6.2	High-Level Control.....	32
2.6.3	Summary	35
2.7	Reinforcement Learning.....	35
2.7.1	Industrial Process Control	37
2.7.2	Algorithm Background	39
2.7.3	Summary	41
3	Hydrocyclone Operational State Detection	42
3.1	Laboratory Hydrocyclone: Single Source	42
3.1.1	Data Collection	42
3.1.2	Model Development	44
3.1.3	Testing and Discussion	46
3.2	Industrial hydrocyclone: Multiple Sources	50
3.2.1	Data Collection	50
3.2.2	Model Development	51
3.2.3	Testing and Discussion	55
3.3	Industrial Hydrocyclone: Single Source.....	58
3.3.1	Data Collection	59
3.3.2	Model Development	61
3.3.3	Testing and Discussion	65
3.3.4	Image-Based Sensor Health Check.....	81
3.4	Summary	90

4	Hydrocyclone Underflow Particle Size Inference.....	91
4.1	Laboratory Hydrocyclone.....	91
4.1.1	Data Collection.....	91
4.1.2	Footage Classification.....	93
4.1.3	Particle Size Regression.....	98
4.2	Industrial Hydrocyclone.....	103
4.2.1	Data Collection.....	103
4.2.2	Image Only Modelling.....	112
4.2.3	Sensor Only Modelling: Hydrocyclone.....	119
4.2.4	Sensor Only Modelling: Extended.....	126
4.2.5	Image and Sensor Modelling.....	136
4.2.6	Discussion.....	142
4.3	Summary.....	148
5	Reinforcement Learning Hydrocyclone Control.....	149
5.1	Experimental Design.....	150
5.1.1	Hydrocyclone Simulation.....	150
5.1.2	Reinforcement Learning Controller.....	151
5.1.3	Simulator-Control Interaction.....	153
5.2	Case studies.....	158
5.2.1	State Representation.....	158
5.2.2	Unsolvable Instances.....	160
5.2.3	Partial Observability Extension.....	164
5.2.4	Set Point Variability.....	170
5.2.5	Noise.....	174
5.3	Discussion.....	184
5.3.1	System.....	184
5.3.2	Actions.....	185
5.3.3	Training.....	187
5.4	Summary.....	188

6	Conclusion	189
6.1	Recommendations	192
	Appendices	197
	Appendix A. Multiple Source Industrial Hydrocyclone Image Datasets Frames Breakdown	197
	Appendix B. Single Source Industrial Cyclone A Image Datasets Frames Breakdown	199
	Appendix C. Three-State Confusion Matrices for Each Cyclone A Test Dataset and Crop Method	200
	Appendix D. Operating and Particle Size Information of Laboratory Experimental Runs	202
	Appendix E. Laboratory Footage Frames Used in Training and Validation Dataset Construction	203
	Appendix F. Laboratory Footage Frames Used in Test Dataset Construction	204
	Appendix G. CycloSim Data_File Spreadsheet Containing Supplied Base Case Data	205
	Appendix H. Copyright Permissions	206
	References	233

List of Abbreviations and Acronyms

ANN	Artificial neural network
BCQ	Batch-constrained deep Q-learning
CIMM	Centro de Investigacion Minera y Metalurgica
CNN	Convolutional neural network
d50	Cut size
d50 _c	Corrected cut size
DDPG	Deep deterministic policy gradient
DPG	Deterministic policy gradient
DQN	Deep Q-network
EIT	Electrical impedance tomography
EROM	Emergency run-of-mine
F	False
fc	Fully connected
FPS	Frames per second
f/	F-number (optics)
GPU	Graphics processing unit
ILSVRC	ImageNet Large Scale Visual Recognition Challenge
LSTM	Long short-term memory
MAE	Mean absolute error
MCC	Matthews correlation coefficient
MDP	Markov decision process
MOSSE	Minimum output sum of squared error
MPC	Model predictive control
MSE	Mean squared error
MV	Manipulated variable
OSM	Oversize monitoring system
PID	Proportional-Integral-Derivative
PPO	Proximal policy optimisation
PV	Process variable
P50	Screen aperture size 50% of solids would pass through

P80	Screen aperture size 80% of solids would pass through
ReLU	Rectified linear unit
RL	Reinforcement learning
RNMPC	Robust non-linear model predictive control
RTO	Real-time optimisation
R ²	Coefficient of determination
SAG	Semi-autogenous grinding
SD	Standard deviation
SP	Set point
SVC	Support vector classification
SVM	Support vector machine
SVR	Support vector regression
T	True
TD3	Twin delayed deep deterministic policy gradient
VGG	Visual Geometry Group
v/v	By volume
WASM	Western Australian School of Mines
w/w	By weight

List of Symbols

a	Action
C	Hydrocyclone feed solids concentration (by weight)
D_A	Apex (spigot) diameter
D_C	Hydrocyclone diameter
D_I	Feed inlet diameter
D_V	Vortex finder diameter
ϵ	Error
E_b	Ball mill power draw
E_s	SAG mill power draw
f	Exponent for $\rho_{(s)}$ effect on d_{50_c}
F	Calibration factor
F_{fan}	Incorrect fan predictions
F_{fault}	Incorrect fault predictions
H	Height, tip of hydrocyclone vortex finder to apex
J	Expected return
k	Class
K_p	Controller gain
L	Loss
M	Feed sizing %passing modification
N	Number of samples in mini-batch
n	Sampled replay buffer transition in mini-batch
η	Fluid viscosity
o	Observation of the environment
P	Hydrocyclone inlet pressure
p	Transition probability
ρ	Hydrocyclone feed slurry density
ρ_s	Feed solids density
Q	Hydrocyclone feed flow rate
Q	Action-value function
R	Target value

R_b	Hydrocyclone underflow (or ball mill) solids flow rate
R_f	Tank solids feed rate
R_s	SAG mill solids feed rate
r	Reward
s	State
T_i	Integral time
T_d	Derivative time
T_{fan}	Correct fan predictions
T_{fault}	Correct fault predictions
T_L	Tank level
t	Time step
τ	Soft-update factor
U_p	Underflow P80 (sensor output)
μ	Policy
V	State-value function
V_F	Feed solids concentration (by volume)
V_O	Overflow solids concentration (by volume)
V_U	Underflow solids concentration (by volume)
V_{apex}	Degree apex tip extends vertically into cropped frame
W	Weights, learnable neural network parameters
X	Operating state (sensor output)
x	Time horizon
x_{apex}	Centre of the apex horizontal coordinate
x_{crop}	Top left horizontal coordinate for crop position
x_{min}	Top left horizontal coordinate for bounding box
\mathcal{Y}	Probability true class = 1 (binary classification)
y_{apex}	Centre of the apex vertical coordinate
y_{crop}	Top left vertical coordinate for crop position
\mathcal{Y}_k	Probability true class = k -th class
y_{min}	Top left vertical coordinate for bounding box
γ	Discount factor

\hat{y}	Network's output (binary classification)
\hat{y}_k	k -th network output

List of Figures

Figure 2-1 Hydrocyclone Parts and Flow Behaviour	5
Figure 2-2 Hydrocyclone Operating States' Underflow Profiles.....	7
Figure 2-3 Convolution Visualisation	22
Figure 2-4 Convolution Calculation Example	23
Figure 2-5 Pooling Layer Dimensional Reduction Example	24
Figure 2-6 Max Pooling Operation Example	24
Figure 2-7 Example of a Network With Two Hidden and One Output Fully Connected Layers	25
Figure 2-8 Output Calculation for a Single Fully Connected Node	26
Figure 2-9 Reinforcement Learning Agent-Environment Interaction.....	36
Figure 3-1 Hydrocyclone Laboratory Experimental Setup.....	43
Figure 3-2 Example Raw and Training Images	45
Figure 3-3 Laboratory Single Source Three-State Classification Network Architecture	46
Figure 3-4 Laboratory Three-State Detector Test Images and Output Classification	48
Figure 3-5 Laboratory Hydrocyclone Unstable Transition Period	49
Figure 3-6 Example Augmented Images	53
Figure 3-7 Industrial Multiple Source Two-State Classification Network Architecture	54
Figure 3-8 Industrial Two-State Detector Test Images and Output Classification	56
Figure 3-9 Cyclone A Underflow Discharge Region Illustration.....	62
Figure 3-10 Industrial Single Source Three-State Classification Network Architecture	65
Figure 3-11 Test Images Crop Position Shifts	66
Figure 3-12 Misclassified Test Images Crop Position Shifts.....	72
Figure 3-13 Region Percentage of Misclassified Crop Position Shift Test Images.....	72
Figure 3-14 Tracker 1 Bounding Box Descriptors During Lens Cleaning.....	83
Figure 3-15 Tracker 2 Bounding Box Descriptors During Lens Cleaning.....	84
Figure 3-16 Tracker 2 Bounding Box Descriptor Difference During Lens Cleaning ...	85
Figure 3-17 Long-Term Bounding Box Descriptor Difference.....	86

Figure 3-18 Long-Term Bounding Box Descriptor Difference With Revised Tracking Region.....	88
Figure 4-1 Hydrocyclone Laboratory Experimental Setup Viewed From Camera Perspective.....	92
Figure 4-2 Image Types and Crop Positions.....	94
Figure 4-3 Hydrocyclone Underflow Footage Classification Network Architecture..	95
Figure 4-4 Hydrocyclone Underflow Particle Size Regression Network Architecture	99
Figure 4-5 Laboratory Validation Dataset Measured and Predicted Underflow P80	100
Figure 4-6 Laboratory Test Dataset Measured and Predicted Underflow P80	102
Figure 4-7 Examples of Illumination Variation Throughout a Single Day	104
Figure 4-8 Experimental Setup at IGO Nova Primary Hydrocyclone Cluster.....	105
Figure 4-9 Banded Footage Example Frames	106
Figure 4-10 IGO Nova Simplified Primary Grinding Circuit Diagram	108
Figure 4-11 Industrial Dataset’s Hydrocyclone Underflow P80 in Order Sampled .	110
Figure 4-12 Industrial Hydrocyclone Underflow P80 Against Sampling Time.....	112
Figure 4-13 Example Video Frame Showing the Crop Location	113
Figure 4-14 Convtune Model Validation Dataset Measured and Predicted Underflow P80.....	115
Figure 4-15 Convtune Model Training Dataset Measured and Predicted Underflow P80.....	116
Figure 4-16 Convtune Model Test NA Dataset Measured and Predicted Underflow P80.....	117
Figure 4-17 Convtune Model Test NB Dataset Measured and Predicted Underflow P80.....	118
Figure 4-18 Neural Network Architecture for Cycsensor Model.....	120
Figure 4-19 Cycsensor Model Validation Dataset Measured and Predicted Underflow P80.....	121
Figure 4-20 Cycsensor Model Training Dataset Measured and Predicted Underflow P80.....	122

Figure 4-21 Cyclesensor Model Test NA Dataset Measured and Predicted Underflow P80.....	123
Figure 4-22 Cyclesensor Model Test NB Dataset Measured and Predicted Underflow P80.....	124
Figure 4-23 Cyclesensor Model Components Contribution to Output Values for Test NA Dataset	125
Figure 4-24 Hydrocyclone Feed Slurry Density Readings for All Datasets.....	126
Figure 4-25 Neural Network Architecture for Gcsensor Model	127
Figure 4-26 Gcsensor Model Validation Dataset Measured and Predicted Underflow P80.....	129
Figure 4-27 Gcsensor Model Training Dataset Measured and Predicted Underflow P80.....	130
Figure 4-28 Gcsensor Model Test NA Dataset Measured and Predicted Underflow P80	131
Figure 4-29 Gcsensor Model Components Contribution to Output Values for Test NA Dataset	132
Figure 4-30 SAG Mill Solids Feed Rate Readings for All Datasets.....	133
Figure 4-31 Gcsensor Model Test NB Dataset Measured and Predicted Underflow P80	134
Figure 4-32 Gcsensor Model Components Contribution to Output Values for Test NB Dataset	135
Figure 4-33 SAG Mill Power Draw Readings for All Datasets	136
Figure 4-34 Network Architecture for Conv-gc Model	137
Figure 4-35 Conv-gc Model Validation Dataset Measured and Predicted Underflow P80.....	139
Figure 4-36 Conv-gc Model Training Dataset Measured and Predicted Underflow P80	140
Figure 4-37 Conv-gc Model Test NA Dataset Measured and Predicted Underflow P80	141
Figure 4-38 Conv-gc Model Test NB Dataset Measured and Predicted Underflow P80	142

Figure 4-39 Box Plot of the Training Dataset Measured Hydrocyclone Underflow P80	146
Figure 4-40 Test Image Examples Exhibiting Over Exposed Regions.....	147
Figure 5-1 Simulated Open Circuit Hydrocyclone System Diagram.....	151
Figure 5-2 Actor and Critic Neural Network Architectures.....	153
Figure 5-3 Count of Test Episodes Terminated in the Given Number of Steps.....	161
Figure 5-4 Initial and Final Process Variable Set Point Differences for Unsolvable Instances Test.....	162
Figure 5-5 Test Episode Tank Solids Feed Rate Against Steps Taken for Episode to Terminate.....	163
Figure 5-6 Process Variables Plot Over Time for the Nominated Test Episodes.....	164
Figure 5-7 Count of Test Episodes Terminating in the Given Number of Steps.....	166
Figure 5-8 Initial and Final Process Variable Set Point Differences for Training Seed Tests	167
Figure 5-9 Initial and Final Process Variable Set Point Differences for Generalisation Tests	169
Figure 5-10 Solvability of Test Episodes Considering Feed Size Modification and Solids Into Tank	170
Figure 5-11 Process Variables and Relevant Set Points During the Set Point Change Test.....	173
Figure 5-12 Manipulated Variables During the Set Point Change Test	174
Figure 5-13 Initial and Final Process Variable Set Point Differences for State 3 Noise A Tests	176
Figure 5-14 Initial and Final Process Variable Set Point Differences for State 3 Noise B Tests	177
Figure 5-15 Unmeasured Variables for Unsolved State 5 Variable Noise Test Episodes	179
Figure 5-16 Initial and Final Process Variable Set Point Differences for State 5 Noise A Tests	180
Figure 5-17 Initial and Final Process Variable Set Point Differences for State 5 Noise B Tests	181

Figure 5-18 Process Variables Plot for State 5 Noise B Tests With Decreasing Tank Level182

Figure 5-19 Manipulated Variables Plot for State 5 Noise B Tests with Decreasing Tank Level183

List of Tables

Table 2-1 Hydrocyclone Variable Influence on Corrected Cut Size	8
Table 3-1 Data Augmentation Settings.....	52
Table 3-2 Modelling Method Performance Comparison.....	58
Table 3-3 Data Augmentation Settings Comparison	63
Table 3-4 Cyclone A Fixed Crop Method Combined Testing Three-State Confusion Matrix.....	67
Table 3-5 Cyclone A Random Crop Method Combined Testing Three-State Confusion Matrix.....	67
Table 3-6 Cyclone A Expanded Crop Method Combined Testing Three-State Confusion Matrix.....	68
Table 3-7 Summary of Crop Method Influence on State Prediction	71
Table 3-8 Cyclone A Fixed Crop Method Combined Testing Two-State Confusion Matrix.....	75
Table 3-9 Cyclone A Random Crop Method Combined Testing Two-State Confusion Matrix.....	75
Table 3-10 Cyclone A Expanded Crop Method Combined Testing Two-State Confusion Matrix.....	76
Table 3-11 Cyclone B Fixed Crop Method Testing Three-State Confusion Matrix....	79
Table 3-12 Cyclone B Fixed Crop Method Testing Two-State Confusion Matrix.....	79
Table 3-13 Cyclone C Fixed Crop Method Testing Three-State Confusion Matrix	80
Table 4-1 Training Dataset Footage Classification Model Accuracy.....	96
Table 4-2 Validation Dataset Footage Classification Model Accuracy	96
Table 4-3 Test Dataset Footage Classification Model Accuracy	97
Table 4-4 Laboratory Validation Dataset Hydrocyclone Underflow P80 Performance Statistics	100
Table 4-5 Laboratory Test Dataset Hydrocyclone Underflow P80 Performance Statistics	101
Table 4-6 Daily Feed Ore Blend.....	111
Table 4-7 ConvTune Model Validation Dataset Hydrocyclone Underflow P80 statistics	115

Table 4-8 Convtune Model Test NA Dataset Hydrocyclone Underflow P80 Statistics	117
Table 4-9 Convtune Model Test NB Dataset Hydrocyclone Underflow P80 Statistics	118
Table 4-10 Cycsensor Model Validation Dataset Hydrocyclone Underflow P80 Statistics	120
Table 4-11 Cycsensor Model Test NA Dataset Hydrocyclone Underflow P80 Statistics	122
Table 4-12 Cycsensor Model Test NB Dataset Hydrocyclone Underflow P80 Statistics	123
Table 4-13 Gcsensor Model Validation Dataset Hydrocyclone Underflow P80 Statistics	128
Table 4-14 Gcsensor Model Test NA Dataset Hydrocyclone Underflow P80 Statistics	131
Table 4-15 Gcsensor Model Test NB Dataset Hydrocyclone Underflow P80 Statistics	134
Table 4-16 Conv-gc Model Validation Dataset Hydrocyclone Underflow P80 Statistics	138
Table 4-17 Conv-gc Model Test NA Dataset Hydrocyclone Underflow P80 Statistics	140
Table 4-18 Conv-gc Model Test NB Dataset Hydrocyclone Underflow P80 Statistics	141
Table 4-19 Trained Models' Mean Absolute Error for All Datasets.....	143
Table 4-20 Test NA Average Hydrocyclone Underflow P80 Difference for Data Collection Event	144
Table 4-21 Test NB Average Hydrocyclone Underflow P80 Difference for Data Collection Event	144
Table 5-1 Example Process Variable Set Point Differences and Corresponding Reward	158
Table 5-2 State Representation and Training Period Test Performance Metrics....	160
Table 5-3 Summary of Hydrocyclone Feed Sizing Information for Modified Base Case	165

Table 5-4 Test Performance Metrics for Policy Trained With Different Random Number Seeds	166
Table 5-5 Summary of Hydrocyclone Feed Sizing Information for Extreme Modified Base Case.....	168
Table 5-6 Test Performance Metrics for Policy Trained With Fixed Set Points.....	171
Table 5-7 Test Performance Metrics for Policy Trained With Varied Set Points.....	172
Table 5-8 Standard Deviation of Random Noise Applied to Measured Variables for Each Test	175
Table 5-9 Test Performance Metrics for Policies Trained With Varied Noise Conditions	177

1 Introduction

1.1 Background

Advances in machine learning through the successful incorporation of deep neural networks have seen impressive results in numerous domains. Human level image classification task performance in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC; Russakovsky et al., 2015) was exceeded by Microsoft's ResNet (He et al., 2015) using a deep convolutional neural network (CNN). Mnih et al.'s (2015) deep Q-network (DQN) incorporated CNNs and deep reinforcement learning (RL) to learn to play a range of Atari 2600 computer games at beyond human-level performance, with only game screen pixels and score as inputs.

The use of CNNs for game board representation along with deep RL and Monte Carlo tree search allowed AlphaGo to beat champion players of the strategy board game Go (Silver et al., 2016). These results in a range of sensing, control, and decision-making tasks have driven interest in how such techniques can be applied to industrial problems. With the availability of open-source software such as Keras (Chollet & others, 2015a), Tensorflow (Abadi et al., 2015), DIGITS (NVIDIA Corporation, 2014) and Caffe (Jia et al., 2014), along with the release of pretrained CNN models, this area of research and development has become more accessible.

Hydrocyclones form a core component in a range of mineral processing circuits, in particular comminution circuits. They influence the particle transfer size to downstream processing units and consequently valuable mineral liberation. Through this their controlled operation has a meaningful effect on operating costs, such as through aiding in efficient comminution and the minimisation of reagent requirements for flotation or leaching circuits, along with ultimately the achievable final product recovery and grade.

Continuous monitoring of hydrocyclones typically only occurs at the cluster level (e.g., through tracking flow rates, slurry density, and feed inlet pressure) rather than individual hydrocyclone performance. By monitoring individual hydrocyclones another level of detail becomes available to assess the circuit, improving the ability to detect and respond to operational faults and product changes. One such proposed

option is image analysis of hydrocyclone underflow discharge appearance, which has been used to determine the operating state (Janse van Vuuren et al., 2011; Janse van Vuuren et al., 2010) and indicated as containing useful information for particle size inference (Aldrich et al., 2015; Uahengo, 2014). Given the performance of CNNs for image analysis tasks, their ability to produce accuracy and reliability image-based hydrocyclone sensors was worth investigating.

Within mineral processing circuits hydrocyclones are also typically operated by low-level control such as Proportional-Integral-Derivative (PID) control loops. These methods can be implemented with the available cluster level sensor readings to aid in providing stable operation, though not necessarily optimal control. High-level advanced control for comminution circuits have been developed incorporating techniques such as model predictive control (MPC; Mintek, 2011b) and real-time optimisation (RTO; Coetzee & Ramonotsi, 2016). Though given the requirement of steady-state modelling for RTO and process dynamic modelling for MPC this can be challenging to achieve on complex and insufficiently defined mineral processing systems (Jiang et al., 2018). For RL a policy to control the system is developed through interacting with it and can occur through model-free methods such as the deep deterministic policy gradient (DDPG) algorithm (Lillicrap et al., 2015). Given DDPG's indicated extension to process control (Spielberg, 2017), deep RL presents an alternative for industrial systems to explore for potential benefits and limitations.

1.2 Objectives

The broad aim of this project was to explore the potential for CNNs to be used in image-based monitoring of hydrocyclones, and RL in the control of hydrocyclones.

There were three core objectives which were investigated

1. The applicability of CNNs for image-based hydrocyclone state detection, through monitoring the underflow discharge. Ideally, the resulting sensor would require little to no site-specific retraining and able to be utilised under industrial conditions.
2. The applicability of CNNs for image-based hydrocyclone underflow particle size inference, through monitoring the underflow discharge. Whether an

industrial sensor can be trained for a specific hydrocyclone from the underflow images alone or through incorporation with other sensor information was examined.

3. The applicability of RL for hydrocyclone control was explored, incorporating the outputs from the sensors developed above (or a related competing method).

While the outcome of each objective is valid as a standalone sensor or control method, their interrelation tied them together in this project as a whole. An image-based state detector would be required for operational fault detection. It would also provide a check that the output of the image-based particle size sensor is likely to be valid, as calibration would only be performed on the desired fan state operation. The introduction of these additional variables opens the possibility of more advanced process control options, of which RL could hold potential for the typically partially defined and noisy circuits containing hydrocyclones.

1.3 Thesis Outline

The thesis structure consists of an introductory first chapter, followed by a literature review in the second chapter. The literature review initially covers the basics of hydrocyclone operation and empirical modelling. The current progress of research and commercialisation of hydrocyclone operating state and particle size monitoring is then covered. Next background information on CNNs is included as they form the basis of the subsequent image-based monitoring investigations. Hydrocyclone control options are then explored, followed by a focus on RL and specifically the DDPG (Lillicrap et al., 2015) algorithm to form the basis of the subsequent hydrocyclone RL control investigation.

The third through fifth chapters will then cover the hydrocyclone monitoring and control investigations undertaken, with methods and methodologies used covered within their relevant section. The third chapter explores the use of CNNs for image-based hydrocyclone state detection. Initially covering laboratory-based proof of concept work, then moving onto more industrially relevant conditions, issues, and applied sensor considerations. The fourth chapter explores the use of CNNs for


image-based hydrocyclone underflow particle size estimation. Again, with initial laboratory-based proof of concept work undertaken, followed by an investigation under industrial conditions and with consideration for including additional sensor information.

The fifth chapter explores the use of RL for hydrocyclone control with a DDPG (Lillicrap et al., 2015) based algorithm interacting with a simulated open circuit hydrocyclone system. The simulated environment was progressively modified to represent challenges and requirements for industrial hydrocyclone control, with the resulting performance assessed. The final sixth chapter concludes the findings of the investigations and presents recommendations for future work.

2 Literature Review

2.1 Hydrocyclone Basics

Hydrocyclones have found widespread use for solids classification and dewatering tasks in mineral processing plants, due to their simplicity and relatively small footprint (Napier-Munn et al., 2005). Figure 2-1 illustrates some of the key physical components of hydrocyclones and the flow behaviour during operation. Slurry enters the hydrocyclone via the feed inlet, where design and physical constraints within the unit result in the production of a spiral flow which travels downwards and exits via the apex (Napier-Munn et al., 2005). This flow causes the particles of the slurry stream to be forced towards the wall of the unit, thus displacing water which moves towards the centre of the unit, with a secondary inner spiral flow forming which travels upwards and exits via the vortex finder (Flintoff & Knorr, 2019; Napier-Munn et al., 2005). Particle separation occurs through the balance of forces involved, with the coarse/dense particles along with minimal water forming the underflow stream. While the effect of fluid drag forces on the fine/light particles results in them drawn with the water to form the overflow stream (Flintoff & Knorr, 2019; Napier-Munn et al., 2005).

Figure 2-1*Hydrocyclone Parts and Flow Behaviour*

Note. Hydrocyclone flow behaviour (left) and key parts (right). From *Mineral Comminution Circuits: Their Operation and Optimisation* (p. 310), by T. J. Napier-Munn, S. Morrell, R. D. Morrison, and T. Kojovic, 2005, Julius Kruttschnitt Mineral Research Centre. Copyright 2005 by Julius Kruttschnitt Mineral Research Centre. Reprinted with permission.

Typical operation for classification purposes sees the generation of an air core within the unit and the characteristic fan shaped underflow profile. Under conditions where the air core collapses the underflow discharge stream takes on a rope-like appearance. For classification purposes this is usually considered a fault state exhibiting a higher underflow solids concentration and coarser overflow discharge (Napier-Munn et al., 2005). A number of studies have been undertaken to predict

conditions where this roping state is likely to occur, with such roping criteria presented in Equation 2-1, Equation 2-2, and Equation 2-3.

Concha et al.'s (1996; as cited in Gupta & Yan, 2006) roping criteria

$$\frac{D_A}{D_V} < 0.45 \quad (2-1)$$

where D_A = apex (spigot) diameter (cm)

D_V = vortex finder diameter (cm)

Mular & Jull's (1980; as cited in Gupta & Yan, 2006) roping criteria

$$V_U > 0.5385 V_O + 0.4911 \quad (2-2)$$

where V_U = underflow solids concentration by volume (v/v)

V_O = overflow solids concentration (v/v)

Laguitton's (1985; as cited in Gupta & Yan, 2006) roping criteria

$$V_U \geq 0.56 + 0.20 (V_F - 0.20) \quad (2-3)$$

where V_F = feed solids concentration (v/v)

If the underflow solids content becomes too high to discharge freely or if foreign materials obstruct the hydrocyclone apex/spigot an underflow blockage can occur. Then with minimal, if any, material able to discharge from the underflow the hydrocyclone feed reports directly to the overflow, with the intended solids classification or dewatering task unable to occur. Figure 2-2 illustrates the key operating states' underflow profiles, with the desired fan shape of normal operation contrasted with the rope-like appearance of the roping state or lack of underflow discharge of the blocked state.

Figure 2-2*Hydrocyclone Operating States' Underflow Profiles*

Note. From “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 2 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Reprinted with permission.

For assessing the solids classification performance of a hydrocyclone, cut size (d_{50}) is a commonly used efficiency parameter being the size of feed particles which would be distributed 50:50 between the discharge streams. Corrected cut size (d_{50c}) is then an adjustment to account for separation deemed to occur by classification forces alone, neglecting particles assumed to bypass classification such as via the water in the underflow stream (Napier-Munn et al., 2005). Other sources also noted to contribute imperfect classification in hydrocyclones include the bypassing of coarse feed material to the overflow, the capture of fine particles in the thickened coarse particle flow, turbulent mixing within the unit, and specific physical properties (e.g., size, shape, and density) of the solids involved (Napier-Munn et al., 2005). Table 2-1 summarises how hydrocyclone design, operating, and feed slurry properties can influence the corrected cut size, and thus how hydrocyclone performance can be manipulated to reach desired product size requirements.

Table 2-1*Hydrocyclone Variable Influence on Corrected Cut Size*

Variable Change	Cut Size (d_{50c})
↑ Apex/Spigot Diameter	Finer
↑ Hydrocyclone Length	Finer
↑ Feed Solids Specific Gravity	Finer
↑ Feed Flow Rate or Pressure	Finer
↑ Vortex Finder Diameter	Coarser
↑ Inlet Diameter	Coarser
↑ Hydrocyclone Diameter	Coarser
↑ Feed Solids Concentration	Coarser
↑ Number of Hydrocyclones Online (with overall feed flow rate constant)	Coarser

Note. A summary of the effect of increasing the stated design, operating and slurry property variables on the corrected cut size, as described in Gupta and Yan (2006, p. 384) and Napier-Munn et al. (2005, p. 329).

2.2 Empirical Models

A number of empirical hydrocyclone models have been developed to predict key engineering variables and find use in design and optimisation work. One of the earliest models was that by Plitt (1976), where equations for inlet pressure, overflow to underflow volumetric split, separation sharpness and corrected cut size were developed by correlating experimental data with a range of hydrocyclone design and operating parameters. Further revision of the Plitt model led to the Plitt-Flintoff model (Flintoff et al., 1987) including additional calibration factors and refinement of the original equations. Equation 2-4 presents the revised corrected cut size model, including a calibration factor, fluid viscosity effect, and compensation for solids density effect (with normalisation given original silica solids based experimental dataset).

Flintoff et al.'s (1987) Plitt-Flintoff corrected cut size model

$$d_{50c} = \frac{F 39.7 D_c^{0.46} D_i^{0.6} D_v^{1.21} \eta^{0.5} e^{0.063 V_F}}{D_A^{0.71} H^{0.38} Q^{0.45} \left(\frac{\rho_s - 1}{1.6}\right)^f} \quad (2-4)$$

where F = calibration factor
 D_c = hydrocyclone diameter (cm)
 D_i = feed inlet diameter (cm)
 η = fluid viscosity (cP)
 V_F = feed solids concentration (% v/v)
 H = hydrocyclone height, tip of vortex finder to apex (cm)
 Q = feed flow rate (l/min)
 ρ_s = feed solids density (g/cm³)
 f = exponent for ρ_s effect on d_{50c}
(e.g., $f = 0.5$ laminar flow, $f = 1$ turbulent flow)

Gutiérrez and Sepúlveda's (1986; as cited in Sepúlveda, 2012) CIMM model is another empirical model of note. This was also derived from a large experimental dataset and includes similar equations to that of Plitt's model, along with an additional slurry short-circuit correlation. The CIMM model finds use in Moly-Cop Tools Excel spreadsheet based open circuit hydrocyclone simulator Cyclosim_Single (Sepúlveda, 2012). The Nageswararao model (Nageswararao, 1995) presents correlations for water and slurry volumetric recovery to underflow, throughput to inlet pressure, corrected cut size, and reduced efficiency curve. With the Nageswararao model and subsequent modifications finding use in JKtech's steady state comminution circuit simulator JKSimMet (Nageswararao, 1995; Napier-Munn et al., 2005). Narasimha et al. (2014) built upon historic hydrocyclone empirical modelling datasets to address limitations that were identified. They incorporated learnings from hydrocyclone computational fluid dynamics studies to produce models, including for sharpness of cut, which have been incorporated in the recent JKSimMet version 6.0.

Subsequent works have extended upon the traditional empirical models with additional variables and modelling techniques. Eren et al. (1996) used a shallow artificial neural network (ANN) to predict hydrocyclone corrected cut size. This was based on closed circuit laboratory hydrocyclone data using silica ore, with the dataset

split 50% training and 50% testing. The resulting model's predictions were comparable with those of a previous empirical model developed from this experimental system (Gupta & Eren, 1990) which used the same input variables: feed flow rate, feed solids concentration, apex diameter, hydrocyclone height, and slurry temperature. The ANN model's predictive performance was also better than that of Plitt's model, which was expected given the ANN was trained on data specific to this system. Eren, Fung, Wong, and Gupta (1997) then extended upon the previous ANN model with additional unconventional variables, such as overflow and underflow flow rates and solids split ratio, yielding further improvements.

Further investigation into the use of shallow ANNs on predicting the corrected cut size produced by a hydrocyclone test rig was undertaken by van Loggenberg et al. (2016). Their dataset was split 60% training, 20% validation, and 20% testing with an increasing number of design and operating variable added to the ANN models trialled. Variables such as inlet pressure, apex diameter, feed solids concentration, feed flow rate and angle of underflow discharge were investigated. The best performing model comprised eight input variables, including unconventional variables also considered by Eren, Fung, Wong, and Gupta (1997) specifically overflow flow rate and density along with underflow flow rate. This model was then compared to the empirical Plitt-Flintoff model, with the ANN based model having a lower prediction error on a range of metrics.

When sufficiently characterised and calibrated, traditional empirical models are powerful tools in mineral process design and optimisation. Recent work with ANNs, and including additionally measured variables, have shown that improved system specific modelling for the likes of corrected cut size can be achieved.

2.3 Operational State Detection

The development of online hydrocyclone state detection has occurred through exploiting different observable phenomenon related to these distinct operating conditions. These investigations have led to a range of sensing options, many of which have progressed from research to commercialisation. This section details several detections methods, focussing on those which could plausibly be

retrofit to existing hydrocyclone systems, and subsequently accelerate their uptake in industry. Thus, methods such as the gravimetric sensor discussed in Neesse et al. (2004), which exploits solids loading differences exhibited by the operating states, was not deemed relevant as the implementation requires flexible feed and overflow piping to facilitate weighing cell function.

2.3.1 Electrical Tomographic Methods

Gutiérrez et al. (2000) investigated using electrical impedance tomography (EIT) to monitor the internal flow behaviour of a hydrocyclone. In a single plane around the external circumference of a hydrocyclone (just below the feed inlet) 16 electrode plates were mounted to facilitate resistivity imaging. A range of calcium silicate slurry feed solids concentrations and flow rates were trialled. When formed, the air core presented as a light region in the grey scale image near the centre of the unit, due to it having a higher resistivity than water present in the slurry. The authors proposed that EIT could thus be used for fault state detection, such as by monitoring the mean and standard deviation of the data, noting that without an air core the slurry within the unit appears more uniform (resulting in a lower standard deviation).

Williams et al. (1999) applied electrical resistance tomography to both pilot plant and operating industrial hydrocyclones. Eight planes of electrodes were distributed along the length of the hydrocyclone (16 electrode discs per plane), allowing 3D conductivity mapping to also be investigated. The presence and nature of the air core was able to be imaged through the conductivity differences of the phases within the hydrocyclone. This allowed for identification of roping, underflow blockage, and spigot failure or excessive wear fault states. Williams et al. (1999) demonstrates the robustness for electrode use in industry and the suitable imaging frequency of electrical tomographic techniques (e.g., in the order of 500 FPS for EIT indicated by Gutiérrez et al., 2000).

Metso Outotec (2021) have released the commercial product CycloneSense where 12 electrodes are contained within a gasket that is then mounted between the hydrocyclone upper and lower cones, enabling it to be readily retrofit to existing

hydrocyclones. This demonstrates the maturity tomographic methods have now reached for continuous hydrocyclone air core monitoring in an industrial setting.

2.3.2 Acoustic Methods

Olson and Waterman (2006) patent for Krebs Engineers (now FLSmidth) presents the use of an ultrasonic sensor in determining whether a hydrocyclone is operating in fan or rope state. The sensor is attached to a cylindrical splash skirt, which is mounted to the hydrocyclone body just underneath the apex. This system exploits the fact that under typical fan operation the underflow will strike the splash skirt and produce a characteristic signal. As the fan collapses the impact with the skirt will change (along with the signal detected), until ultimately under rope conditions the underflow does not impact the splash skirt. Buttler et al. (2019), a subsequent FLSmidth patent application, refers to a wireless sensor extension with the commercially available FLSmidth KREBS SmartCyclone wireless sensor having a 1 s measurement rate and 5 s reading update rate (FLSmidth, 2021).

Nesse et al. (2004) demonstrated that the differences in vibration exhibited by hydrocyclones when operating in fan or rope state can be exploited for state detection. In their case, a single accelerometer sensor was mounted externally on the hydrocyclone body near the apex, with amplification and low pass filtering applied in produce a vibration frequency spectrum. Bowers et al.'s (2019) patent for Emerson Electric details a system using dual accelerometer sensors, attached to each hydrocyclone near the overflow and underflow regions. Band pass filtering is applied to the vibration signals, and by comparing the signals over time and in relation to each other in certain frequency ranges roping and blocked underflow operating states can be identified.

Putz De La Fuente's (2019) patent application also presents a vibration-based system, though in this case requiring a modified hydrocyclone overflow pipe setup. The sensor assembly is attached to an opening in the overflow pipe, at the hydrocyclone end of the pipe. The assembly consists of a membrane which forms part of the pipe's internal wall, and an accelerometer sensor attached to the external

side of the membrane. By monitoring the membrane's vibration characteristics in certain frequency ranges roping operation can be identified.

An acoustic sensor which clamps around the outside of each individual hydrocyclone overflow pipe is used in CiDRA's CYCLONetrac Oversize Monitoring (OSM) system (CiDRA, 2018). The impact of large particles (≥ 6 mm) on the walls of the pipe creates strain which is detected by the sensor encompassing the pipe (O'Keefe et al., 2014), and with a 4 s reading update rate (CiDRA, 2018). O'Keefe et al. (2014) demonstrated that by monitoring the frequency and history of large particle impacts, events in which excessive coarse particles are present in the overflow can be identified (e.g., hydrocyclone roping or blockage). Though they also note that conditions such as high hydrocyclone feed solids concentration and damage to mill trommel screens or hydrocyclones can also result in excessive coarse material in the overflow stream, so interpretation of the readings based on individual hydrocyclones and the circuit as a whole is required. For hydrocyclone circuits in which the feed stream contains particles greater than the detectable size, CYCLONetrac OSM presents an option for fault state detection that can be easily retrofitted.

2.3.3 Underflow Probe Methods

Hulbert's (1992) patent for Mintek details the use of a probe suspended from an arm which can rotate about a fixed point. They describe the probe as being encouraged to ride on the external surface of the underflow stream, by gravity or an applied spring. The resulting angle of the arm could then be detected by optical or ultrasonic based methods, and subsequently allow the underflow discharge profile to be determined and correlated to operating state. Strudwicke et al.'s (2017) patent application for Weir Minerals details an extended probe, fixed to the spigot housing, which protrudes into the hydrocyclone underflow discharge region. A vibration sensor, whether mounted within a channel in the probe or closer to the mounting point, could then detect the vibrations caused by the underflow discharge fan striking the probe. By positioning the probe in such a way that a roping discharge stream does not contact it then the absence of vibration could be used to signify roping state (or potentially blockage).

While both methods describe the requirement for abrasion resistant materials, the nature of these physical contact-based methods still present the potential for wear, damage, and material build-up to influence their long-term viability compared to some of the non-contact state detection methods that have also been proposed. Neesse et al. (2004) present non-contact methods in which a capacitance probe, or laser proximity switch, in the underflow discharge region is used to detect whether a fan discharge is occurring or not. While these would also be protected and not intended to be in direct contact with the underflow discharge, the probe's presence in the underflow region still increases the potential for damage compared to remote sensing methods, such as optical.

2.3.4 *Optical Methods*

Section 2.3.4 contains material previously published in "Operational state detection in hydrocyclones with convolutional neural networks and transfer learning," by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 2 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier.

Optical sensing methods for hydrocyclones exploit the fact that the underflow discharge profile acts as an external means of identifying the operating state, as noted in Section 2.1 and illustrated in Figure 2-2. Neesse et al. (2004) describes how a charge-coupled device camera could be used to capture the reflection of a 2D laser off the underflow discharge. With the laser beam presented at an angle and the camera perpendicular to the hydrocyclone, they show the collapse of a broad fan shaped underflow would see the reflected projection line both shift lower and become narrower. Pattern recognition could then be used in assessing the profile and subsequently classifying the discharge as fan or rope, with an extension for a lack of significant projection line to represent underflow blockage being plausible. Similarly, Mintek (2011a) presents CyLas as a commercially available laser-based sensor, which monitors the distance from mounting point to the underflow discharge to determine discharge angle (Giglia & Aldrich, 2020).

As previously discussed in Giglia and Aldrich (2020), another optical method proposed to exploit the physical differences in the underflow profile of different

operating states is image analysis, thus negating the use of lasers and providing a richer source of information. Early work by Petersen et al. (1996) and van Deventer et al. (2003) was based on the use of a video camera and additional lighting to determine the underflow spray angle on both laboratory and industrial hydrocyclones. Video frames were enhanced via filtering and thresholding techniques to isolate the edges of the spray profile from which the spray angle could be calculated. Janse van Vuuren et al. (2010) and Janse van Vuuren et al. (2011) extended this work, but with a slightly different approach. Rather than processing the video frames as a whole, only the underflow width at a fixed horizontal plane was considered, thus tracking the same region over time given by a fixed camera. Individual frames were extracted from the video footage, converted to grey scale, and enhanced via contrast/brightness adjustment to improve the prominence of the underflow stream. The pixels' intensities along a horizontal frame were then extracted to give 1D intensity spectrum from which the spray profile could be characterised (Giglia & Aldrich, 2020).

In Janse van Vuuren et al.'s (2011) case, for every 1,800 frames the spectra are assessed to establish a search region from which motion detection between consecutive frames was used to determine the underflow width. The phase space was constructed using the 30 frames underflow width moving average for the current frame and a specified historic frame (to incorporate system dynamics). A single class support vector machine (SVM) classifier, using the slurry fan state as training data, and tested against the fault states (dilute fan or roping) was trialled on a number of different ore types. They found this method to be robust in handling system vibrations, background, and random noise, but had issues with foreground noise and conditions with low contrast between discharge and background.

For Janse van Vuuren et al. (2010) a technique involving image normalisation and the quantile distribution was used to determine which points to include in the underflow width measurement. A range of classification modelling techniques were trialled incorporating underflow width, hydrocyclone feed pressure and solids concentration to differentiate between fan, rope, and unstable operational states (with the potential for extension to detect blocked states also noted). The neural

network based classifier resulted in the highest test set accuracy, though with the transition state detection accuracy distinctly lower than that obtained for the fan and rope states. Overall, the image-based hydrocyclone state detection research covered has demonstrated the potential viability of such methods through the extraction of engineered features along with applying various image processing techniques to aid in discerning the underflow profile (Giglia & Aldrich, 2020).

2.3.5 Summary

Numerous hydrocyclone operational state sensing methods have been presented, including potentially industrially viable or already commercially available options. The gravimetric methods would likely require too extensive modifications to existing systems to warrant retrofitting compared to the other methods presented. The recent release of a commercial tomographic sensor capable of being retrofitted to hydrocyclones shows the maturity of this method for industrial use. The potential service life and maintenance requirements of underflow probe-based methods may limit their appeal compared to remote sensing options. A range of acoustic-based methods are also mature hydrocyclone operational state detection options, with a number of commercial products able to be retrofitted to existing systems. Optical methods do however present a potentially viable alternative state detection method, requiring no modification to the hydrocyclone structure to facilitate their use. Commercial laser-based methods are currently available, but the progress of image-based methods does offer an attractive alternative as a potential non-contact sensing option.

2.4 Online Particle Size Monitoring

Particle size distribution is useful in hydrocyclone performance monitoring, be it that contained in the overflow or underflow stream. Particle size measurement or inference will often take the form of stating the screen aperture a given percentage of solids would pass through (e.g., 80% for P80 and 50% for P50), or the percentage of solids that would pass through a given screen aperture. While manual checking of hydrocyclone stream sizing by screening methods is often performed during process monitoring, methods for automated and sufficiently frequent online

particle size measurement or estimation would be a requirement for continuous control and optimisation.

2.4.1 Sub-sample Measurement

Commercial online sizing measurement of slurry streams are currently available where a sub-sampling cut of the stream is taken and transferred to a unit for analysis. Outotec's PSI 500i analyser uses laser diffraction to provide a volumetric particle size distribution, being applicable to particle sizes 0.5–1000 μm and with a 180 s sample analysis frequency (Metso Outotec, 2020). According to Outotec (2009), the PSI 300 analyser uses a physical calliper which randomly captures and measures particles at a rate of 2 measurements per second. Over a period of time this process can be considered a representative sample and used to present a particle size distribution, being applicable to particle sizes 25–1000 μm and with a 70–90 s sample analysis frequency (Outotec, 2009). The nature of the sub-sampling process for these systems usually results in their use on composite streams rather than the product of individual hydrocyclone.

2.4.2 Soft Sensors

Investigations have also been undertaken to produce particle size predicting soft sensors from available industrial sensor information. Eren, Fung and Wong (1997) used ANNs with operational and design variables (including feed particle size distribution) from a laboratory hydrocyclone system to predict densities and particle size distributions in the overflow and underflow streams. Whilst the resulting model was not specifically used for online particle size predictions, it gives a demonstration that the ANN based empirical cut size model extensions can also be applied to particle size distributions. Zhang and Liang (2016) used available solids and water flow rates, solids concentrations, and mill power data from an operating primary grinding ball mill in closed circuit with a hydrocyclone to predict particle size. The dataset was split 50% training and 50% testing with the six input variables used for both support vector regression (SVR) and a shallow ANN modelling, finding SVR produced a slightly more accurate model.

A commercialised example of a particle sizing soft sensor is part of Mintek's Millstar advanced control system. The Millstar particle size estimator is used to predict hydrocyclone overflow particle size, being based on an empirical model produced following plant test work (Mintek, 2011b). The existing soft sensor research shows that there are variables in hydrocyclone circuits that can be measured and used in performance monitoring. Emerson Electric have worked towards using vibration readings in hydrocyclone monitoring, via sensors mounted externally on the unit near the apex and overflow (Bowers et al., 2019). They have indicated that the incorporation of vibration data along with that of other measured variables can be used in a soft sensor to infer particle size in the overflow and underflow streams (Cahill, 2021).

2.4.3 Overflow Discharge

CiDRA's CYCLONEtrac range of sensors are acoustic-based and used for individual hydrocyclone overflow monitoring. While the CYCLONEtrac OSM system discussed in Section 2.3.2 detects the presence of coarse particles in the overflow stream, it is used mainly for fault or undesirable condition detection. Alternatively, Maron et al. (2018) details CiDRA's CYCLONEtrac particle size tracking system which uses a probe inserted into the discharge end of the hydrocyclone overflow pipe. This sensor converts the stress waves caused by particle impact on the probe into an electrical signal, which is correlated with a mass of material passing or retained by a screen size of interest in the calibration process (Maron et al., 2018). The system is applicable to particle sizes $\geq 75 \mu\text{m}$, with a 4 s reading update frequency (CiDRA, 2019), and has an approximately 18 months probe life due to abrasion from slurry flow (Maron et al., 2018). While the CYCLONEtrac sensors are useful methods for overflow particle size characterisation, they are not currently applied to individual underflow streams. They have been designed to be fitted to piping downstream from the hydrocyclone, while hydrocyclone underflow slurry usually discharges directly from the unit into a communal launder.

2.4.4 Underflow Discharge

As evident by its inclusion in van Loggenberg et al.'s (2016) ANN empirical model extension investigation, and the extreme example of fan versus rope

discharge appearance, underflow spray angle is an optically detectable external indicator of hydrocyclone performance. Hulbert's (1993) patent for Mintek extends upon the physical underflow probe method of Hulbert (1992), discussed in Section 2.3.3, to describe how the measured spray angle could be incorporated into a model along with feed flow rate and solids concentration measurements to infer particle size information of the discharge streams. Mintek's CyLas (Mintek, 2011a) laser-based sensor provides continuous hydrocyclone underflow spray angle monitoring, with an indicated measurement range of 0°–30°. Thus, along with its use for detecting undesirable operating conditions, as noted in Section 2.3.4, the signal could be used in hydrocyclone performance monitoring.

As an alternative optical method, preliminary work was performed by Petersen et al. (1996) for monitoring the underflow spray angle from laboratory and industrial hydrocyclone video footage. A low-pass Gaussian filter and thresholding preprocessing was applied to the grey scale footage frames to aid in discerning the underflow stream from the background. This resulted in a binary image with most background values as zero and underflow stream values as one. Edge detection was then applied to the binary image, with the resulting spray angle able to be correlated with operating parameters (including percentage of <75 µm particles in overflow), though became less reliable during periods of operational instability.

Dubey et al. (2016) analysed images of a laboratory hydrocyclone with a silica feed under various hydrocyclone geometry and operating conditions. Image analysis took the form of RGB to grey scale image conversion, thresholding to produce a binary image, then edge detection and spray angle determination. At each test condition, 10 images were taken (1 FPS) with the spray angle deemed to have an acceptable degree of stability and reproducibility under a given set of conditions. In their study an empirical model was generated for the system, correlating spray angle to several key parameters.

Aldrich et al. (2015) took an alternative approach in analysing the underflow images of a laboratory hydrocyclone by investigating if textural features are usable for particle size classification. Experimental runs from three ore sources and variable solids concentrations were performed, and once stabilised at each interval underflow

images and samples for sizing were taken, leading to a total of 15 runs. Image preprocessing via RGB to grey scale conversion and image normalisation was performed, with a steerable pyramid algorithm and the subsequent decomposition used as feature extraction. The dataset was divided into three classes based on underflow particle size (fine, intermediate, coarse), and a linear discriminant classifier trained. Cross validation was used during model training with each class split randomly into 75% training and 25% testing, with steerable pyramid hyperparameters and inclusion of principal component analysis for feature reduction explored. The average classification model performance over 10 different training/tests split repetitions was approximately 92% accuracy.

Uahengo (2014) undertook underflow mean particle size modelling using the same 300 example laboratory hydrocyclone experimental dataset as (Aldrich et al., 2015), split 80% training and 20% testing. The image features used as inputs were underflow width (extracted in a similar manner to Janse van Vuuren et al., 2011), five grey level co-occurrence matrix statistics, and pixel intensity standard deviation. Modelling was trialled using both multiple linear regression and an ANN using a single hidden layer and 3-fold cross validation on the training set. Ultimately both models were deemed unsuccessful in correlating these image features to mean particle size, though the ANN model showed a better fit to the data. Some of the recommendations Uahengo (2014) proposed to improve upon this work was through the collection of a large dataset, the incorporation of additional measurable parameters into the model, attempting to have the image capture whilst simultaneously performing underflow sampling, and the capture of sequential images for further study, all of which were considered in the experimental design in Section 4.

2.4.5 Summary

A range of online particle size sensing options for hydrocyclone circuits have been presented, with several commercial options available for different purposes. Physical measurement systems are not practical for individual hydrocyclones and have specific particle size and sample frequency requirements. The soft sensor research demonstrates that available design and sensor data can be used to rapidly

infer particle size related information from a hydrocyclone circuit. Emerson Electric's incorporation of vibration data indicating that individual hydrocyclone stream particle size inference is possible through such methods (Cahill, 2021). CiDRA's CYCLONEtrac acoustic sensor presents a rapid method for determining individual hydrocyclone overflow sizing information (CiDRA, 2019), but does not appear applicable to the underflow stream. The profile of a hydrocyclone's underflow has been found to contain optically detectable information that can be correlated to particle size, both in terms of spray angle and textural features. Image-based research to date has used a range of preprocessing, image analysis and machine learning techniques to investigate this relationship.

2.5 Convolutional Neural Networks

The hydrocyclone underflow image-based studies discussed have applied traditional machine learning techniques with engineered features, however recent progress in image analysis has seen improved performance using deep CNNs. Krizhevsky et al.'s (2012) use of CNNs in AlexNet for the 2012 ILSVRC (Russakovsky et al., 2015) resulted in a significant improvement in image classification performance compared to previous computer vision techniques. The ILSVRC makes use of a subset of the ImageNet (Deng et al., 2009) dataset and contains a diverse range of labelled images from 1000 different classes (including some specific subclass examples like dog breeds). The resulting classification database for the ILSVRC contains 1.2 million training images (732–1300 images per class), along with 50 validation and 100 test images per class, creating a large benchmark database for algorithm development and evaluation (Russakovsky et al., 2015). Further improvements in network architectures and implementation techniques ultimately led to Microsoft's ResNet (He et al., 2015) entry being the first CNN to exceed human level performance (5.1% top-5 classification error; Russakovsky et al., 2015) in the 2015 ILSVRC.

2.5.1 Background


There are a range of CNN architecture designs that have found success in image analysis tasks, such as those used in the ILSVRC, and share many of the same core components. The input images to CNNs take the form of a 3D array of pixel intensities with dimensions of the image width by height and either 1 channel depth for grey

scale or 3 channel depth for a RGB images, and may be preprocessed by methods such as rescaling, centring or normalisation.


Convolutional layers use filters (stacked 2D kernels) of a defined size which travel over the input space. The kernels consist of learnable weights (w) which are locally connected to the area of input space they assess and provide a trainable way of extracting meaningful features (Karpathy et al., 2015). Figure 2-3 illustrates the convolution operation between a kernel and region of the input, with the element-wise multiplications summed to produce the output (shown with no bias added) as shown in Figure 2-4. This process is repeated over the entire input space, with the resulting output from a single filter referred to as a feature map. Convolutional layers are then progressively built up to identify higher-level features (e.g., parts and objects) composed of lower-level features (e.g., edges and motifs; LeCun et al., 2015).

Figure 2-3

Convolution Visualisation




Note. From *Figure3*, by A. Naidu, 2019a, GitHub (<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure3.jpg>). Copyright 2019 by Ashwin Naidu. Reprinted with permission.

Figure 2-4*Convolution Calculation Example*


Note. From *Figure4*, by A. Naidu, 2019b, GitHub (<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure4.jpg>). Copyright 2019 by Ashwin Naidu. Reprinted with permission.

Pooling layers are used for dimensional reduction of feature maps, as illustrated in Figure 2-5, where set sized filter travels over each feature map and summarises the key information in that region. Examples include average pooling in which the average activation of the input region is the output, or max pooling where the maximum activation of the input region is the output (Karpthy et al., 2015), as illustrated in Figure 2-6. Dimensionality reduction, or conversely expansion, in terms of depth (channels) can be performed using 1x1 convolution, referred to as cross channel parametric pooling in Lin et al. (2013).

Figure 2-5*Pooling Layer Dimensional Reduction Example*

Note. From *Pool*, by A. Karpathy, 2015c, GitHub (<https://github.com/cs231n/cs231n.github.io/blob/master/assets/cnn/pool.jpeg>).

Copyright 2015 by Andrej Karpathy. Reprinted under an MIT license.


Figure 2-6*Max Pooling Operation Example*

Note. From *Maxpool*, by A. Karpathy, 2015a, GitHub (<https://github.com/cs231n/cs231n.github.io/blob/master/assets/cnn/maxpool.jpeg>). Copyright 2015 by Andrej Karpathy. Reprinted under an MIT license.

Fully connected (fc) layers can find use within the CNN architecture as hidden and output layers which perform a global combination of features from their input (Chollet, 2018). Each node in these layers is connected to every node of the previous layer by learnable weights, as shown in Figure 2-7. Figure 2-8 illustrates the calculation of a single fully connected node output (shown with a bias added). The final output layer for CNNs will have a structure relevant to the task. For multi-class classification there would be a node for each class (k) with a softmax activation, representing a probability distribution over the classes by rescaling all output activations to within the range 0–1 and such that they sum to one (Goodfellow et al., 2016).

Figure 2-7

Example of a Network With Two Hidden and One Output Fully Connected Layers




Note. From *Neural_net2*, by A. Karpathy, 2015b, GitHub (https://github.com/cs231n/cs231n.github.io/blob/master/assets/nn1/neural_net2.jpeg). Copyright 2015 by Andrej Karpathy. Reprinted under an MIT license.

Figure 2-8

Output Calculation for a Single Fully Connected Node

$$n = aw + bx + cy + dz + \text{Bias}$$



Note. Single fully connected node (green font) with four inputs (purple font), learnable weights (orange font) and bias (yellow font).

Deep CNNs are trained through supervised learning in which the network's output is compared to the known values of the training images. Continuing the multi-class classification case, cross entropy loss (L) would compare the true class probability distribution to the network's output probability distribution for an image (Goodfellow et al., 2016), as in Equation 2-5. The loss is used in backpropagation, determining the gradient with respect to the learnable weights, which then guide the changes in network weights to reduce the error through stochastic gradient descent methods (LeCun et al., 2015). As succinctly stated by LeCun et al. (2015) the backpropagation method "is nothing more than a practical application of the chain rule for derivatives" (p. 438).

Multi-class cross entropy loss

$$L(w) = - \sum_k y_k \log \hat{y}_k(w) \quad (2-5)$$

where y_k is probability true class = k -th class

(i.e., 1 for labelled true class, 0 for all other k classes)

\hat{y}_k is the k -th network output

(Murphy, 2012)

A useful property exhibited by CNNs is the potential transferability of learnt features, with Yosinski et al. (2014) exploring this behaviour using an AlexNet (Krizhevsky et al., 2012) implementation and the ImageNet (Deng et al., 2009) dataset. Experiments were performed in which the dataset's classes were split randomly as well as into natural and artificial object classes, to gauge the learning effect on dissimilar tasks. Transfer learning was performed both in terms of using earlier network layers with frozen pretrained weights and only training the top layers with randomly initialised weights, along with fine-tuning in which the pretrained weights from layers earlier in the network can be trained further. They found that both the specialisation of later layers in the network along with co-adaptation between layers affects the ultimate performance when applying transfer learning. While they also noted that the performance of transfer learning diminished with dissimilar tasks, it still held an advantage over training from randomly initialised weights alone. The use of pretrained CNNs also isn't limited to training entirely neural network based classifiers, with Razavian et al. (2014) using a pretrained CNN as a feature extractor (with frozen weights) and then training a SVM classifier on the CNN's fixed sized output feature vector.

2.5.2 Notable Networks

This section will briefly detail several notable CNNs and some of their key features, many of these where top performing entrants in past ILSVRC events. For further details on each network's architectural structure the relevant cited papers along with their implementations on a range of open-source software can be examined.

As introduced in Section 2.5, AlexNet's 2012 ILSVRC performance led to significant interest in CNNs for computer vision. Krizhevsky et al.'s (2012) AlexNet consists of five convolutional layers, three max pooling layers, two fully connected hidden layers and a fully connected output layer. To speed up training beneficial design features include using parallel graphics processing units (GPUs) and the non-saturating rectified linear unit (ReLU) activation function to all convolutional and fully connected layer outputs. To combat overfitting Krizhevsky et al. (2012) performed image augmentation (including random cropping, horizontal flipping, and channel

intensity shifts) which effectively expand the training dataset. They also used dropout with a 50% probability after each fully connected hidden layer. Srivastava et al. (2014) presents dropout as being a regularisation technique by which nodes and their associated weights in the designated layer are randomly set to zero with a given probability during the training process. They propose this technique breaks co-adaptation as no single node can be relied upon to be present during training, and also acts as a form of simplified model averaging, ultimately leading to improved generalisation performance.

Simonyan & Zisserman's (2014) Visual Geometry Group (VGG) entries into the 2014 ILSVRC resulted in the highest localisation and second highest classification performances. The VGG networks focused on depth and a linear structure, with VGG-16 and VGG-19 containing 13 and 16 convolutional layers respectively (with small 3x3 receptive fields), along with five max pooling layers, two fully connected hidden layers and a fully connected output layer. The top classification task performance for the 2014 ILSVRC came from Szegedy et al.'s (2015) GoogLeNet, which makes use of what they called inception modules to produce a wider and deeper network structure. Their inception modules aim to perform multiscale feature extraction on the input, containing 1x1, 3x3, and 5x5 convolution filters along with 3x3 max pooling, and with the output feature maps concatenated. GoogLeNet's architecture makes use of varying arrangements of convolution, pooling, and inception modules to create a network with 21 trainable hidden layers and four max pooling layers. This is then followed by global average pooling of the final feature maps and a fully connected output layer.

He et al.'s (2015) ResNet submission became the first to exceed human level classification performance on the ImageNet dataset at the 2015 ILSVRC. Their residual networks incorporated shortcut connections which facilitate input information transfer to designated subsequent convolutions layers, to aid in combatting performance degradation of very deep networks. The ResNet architectures have a modular structure that can be built upon; all starting with a 7x7 convolutional layer, followed by a 3x3 max pooling layer, then ultimately ending with global average pooling of the final feature maps and a fully connected output layer.

Following the max pooling layer, a number of convolutional layer blocks with varying receptive fields and a short cut connection are used. For example, ResNet50 has 16 blocks (48 convolutional layers) while ResNet152 has 50 blocks (150 convolutional layers), as each block consisting of three convolutional layers (1x1, 3x3, 1x1). After each convolution the ResNet architectures also incorporate Batch Normalization which normalises the layers' output for each training mini-batch, stabilising and speeding up training along with providing a regularization benefit (Ioffe & Szegedy, 2015).

While the CNN architectures discussed have focused on the progressive improvement in accuracy with respect to the ImageNet (Deng et al., 2009) dataset, there is another factor to consider and that is computational efficiency. If these models are to be applied in the continuous analysis of video footage, then being able to perform inference in real time on available hardware would be ideal. One such example of an efficiency focussed CNN architecture is Howard et al.'s (2017) MobileNet, which uses depthwise separable convolutions. The depthwise separable convolution blocks in MobileNet performs depthwise convolution with one 3x3 filter per input channel and then pointwise convolution (1x1 convolution) to combine the previous layer's outputs, which leads to less computational cost than standard 3x3 convolution. MobileNet's structure starts with a 3x3 convolutional layer, then 13 depthwise separable convolution blocks (26 convolutional layers), global average pooling of the final feature maps and a fully connected output layer.

2.5.3 Mineral Processing

In terms of industrially relevant CNN research, Fu and Aldrich (2018) investigated the use of transfer learning on mineral processing related data. In this case, images of flotation froth from arsenic sulphide ore laboratory batch flotation and platinum metal group industrial flotation circuit. Their investigation assessed the performance of AlexNet (Krizhevsky et al., 2012) for feature extraction compared to alternative methods including grey level co-variance matrices, local binary patterns, steerable pyramids, and wavelets. Coupled with random forest modelling, they found the AlexNet transfer learning feature extraction method to be superior to the alternatives examined in terms of coefficient of determination (R^2) for arsenic

concentration in froth estimation and test set error for industrial operating condition classification. Given these results, Fu and Aldrich (2019) then extended upon this work through re-examining the industrial flotation images with a number of pretrained CNNs and applying varying degrees of further training. The fine-tuned ResNet34 (He et al., 2015) implementation resulted in the best test set accuracy, outperforming AlexNet (Krizhevsky et al., 2012) and VGG-16 (Simonyan & Zisserman, 2014) feature extraction networks, and fine-tuned VGG-16.

Further CNN based flotation monitoring investigations have been trialled incorporating historic information in producing froth grade estimates. Zhang et al. (2020) trained a CNN model to produce an initial froth grade and impurity type predictions, using GoogLeNet (Szegedy et al., 2015) for feature extraction, on industrial lead-zinc froth images. Information from key similar historical results stored in memory were then interrogated to refine the predictions. Using several statistical metrics, their method was found to be outperform alternative vision models including simpler CNN models, and an existing fuzzy expert system sensor.

During the period in which this research project was being finalised, our wider research group also published on the use of CNNs for hydrocyclone underflow particle size estimation. Olivier and Aldrich (2021) revisited the 300 example images and underflow particle sizing data from Uahengo (2014) and Aldrich et al. (2015). The dataset was randomly split 70% for training and 30% for validation. Pretrained GoogLeNet (Szegedy et al., 2015) was used up to the final pooling layer, with a new 100 node fully connected layer with ReLU activation followed by a single node fully connected layer added. Both training of the new fully connected layers and fine-tuning the last two inception modules was undertaken until the validation mean squared error (MSE) loss did not improve for 10 epochs, leading to validation set R^2 of 0.851 and 0.910, respectively.

Given the small dataset (limited experimental runs and associated images) used by Olivier and Aldrich (2021) an independent test dataset was not available, so the model's ability to generalise could not be tested. This limits what can be interpreted from the model's performance beyond the indicated fit to the dataset, and improvement yielded by task specific convolutional layer fine-tuning. Collecting

a large dataset, given the potential to overfit to training data, along with consideration for an independent test dataset were key factors in the experimental design in Section 4.

2.5.4 Summary

Through the utilization of technology advancement and large datasets, CNNs have been able to facilitate advances in a range of computer vision tasks. Continued research into network architectures and implementation techniques have produced a range of high performing image classifiers, many of which have been released open source and with weights pretrained on exist datasets (such as ImageNet; Deng et al., 2009). The ability to apply transfer learning techniques to use pretrained CNNs on smaller datasets and even dissimilar tasks, further expands the potential use cases for these networks as highlighted by the mineral processing related research.

2.6 Hydrocyclone Control

A key method for getting the most out of any established circuit in a modern mineral processing plant is through its process control systems. Process control holds a hierarchical structure, with low-level regulatory control for stability and disturbance handling of simpler/predictable systems, high-level advanced regulatory control for multivariable and complex systems, through to advanced control for optimisation (Yench et al., 2015).

2.6.1 Low-Level Control

Wei and Craig's (2009) grinding circuit control survey noted a majority of respondents used proportional–integral–derivative (PID) control in their circuits. They attributed this prevalence to the control method being easily understood, the relatively lower and more widely available skill level to implement than more advanced methods, and a reluctance to employ external vendors for the likes of remote process monitoring. The difference between the control set point (SP) and process variable (PV) is the error (ϵ), such that $\epsilon = SP - PV$. The PID feedback control algorithm considers this difference (proportional), its history (integral), and rate of change (derivative) to determine the controller output, manipulated variable (MV; Wade, 2017). The complete PID algorithm has an ideal form as shown in Equation

2-6, though in practice only some components may be used (e.g., P-only or PI) along with optional modifications to improve performance (Wade, 2017). Further regulatory methods can also incorporate controllers such as feed forward, cascade, override, and dead-time compensation.

Ideal PID control algorithm

$$MV = K_p \left(\varepsilon + \frac{1}{T_i} \int \varepsilon dt + T_d \frac{d\varepsilon}{dt} \right) \quad (2-6)$$

where K_p = controller gain

T_i = integral time

T_d = derivative time (Wade, 2017)

In terms of control strategy, Flintoff and Knorr (2019) noted that for open circuit hydrocyclone systems feed density is the main practically controllable driver of overflow particle size. The feed pressure would then be targeted, through the number of operating hydrocyclones and factoring in feed flow rate, to maintain the pressure in a stable range without requiring frequent changes. Orway Mineral Consultants (2017) discuss their experience and issues encountered around closed circuit hydrocyclone control philosophies, given the complexity the recirculating load brings. They note a preference for control based on feed flow rate adjustment to a pressure target and feed density adjustment to maintaining tank level, as it allows the hydrocyclone cut point to drift to match milling capability. With this method they suggest, overflow density would then be used to guide the number of hydrocyclone in operation. The range of control philosophies they describe as having encountered in grinding circuit operation highlights the subjectivity in the industry for what is the most long-term stable and optimal control behaviour.

2.6.2 High-Level Control

Advanced control methods become relevant for systems with a greater level of process complexity, disturbances, and potential constraints (Wade, 2017). Model-based control, as the name suggests, uses a model of the process constructed from historic data in determining optimal actions whilst considering constraints. Their

performance, however, is dependent on the current strength and validity of the process model the controller is based on (Bascur, 2019). Expert systems attempt to incorporate operator or expert knowledge into rule-based control. With an if-then statement structure and using crisp (true/false) or fuzzy (degree of truth) logic, decisions for the control strategy (such as SP changes) can be developed in a human readable way (Bascur, 2019; Yench et al., 2015). Ultimately, if a model of the system can be produced then model-based methods are suitable, conversely rule-based methods form a model-free alternative (analogous to modelling the operator's behaviour; Bascur, 2019).

Research into advanced methods of hydrocyclone control has been approached in various ways. Gupta and Eren (1990) demonstrated an iterative controller for a laboratory hydrocyclone system with a fixed feed particle size distribution. A model for d_{50c} was produced from experimental data and available sensor information, then used to sequentially adjust three MVs until the d_{50c} SP or all constraints were reached. Wong et al. (2004) extended upon this work using the same system as Gupta and Eren (1990) but investigated using hybrid fuzzy modelling. Fuzzy rules were generated using the Improved Sugeno–Yasukawa modelling method (Tikk et al., 2002) on input-output experimental data from the laboratory hydrocyclone system. Wong et al. (2004) then used memetic algorithms as local search optimisation to refine fuzzy parameters of the identified rules. This produced three fuzzy rules to incorporate into the control, thus resulting in a simple, human readable control strategy.

Stange (1992) presented model-based control interacting with a simulated open circuit desliming hydrocyclone system, with limited measured variables. Using system response data for varying feed size distributions, but maintaining target underflow sizing, a model was constructed using an ANN. The model mapped hydrocyclone feed and overflow density to an operating pressure, to be used for SP updating. This could then be applied to the system without online particle size measurement, where PID controllers were used for tank level and operating pressure control. Model training with varying degrees of noisy data was noted to yield a stable

control response, but increased noise lowered the ability to compensate for the unmeasured feed particle sizing.

Mintek (2011b) presents Millstar as a commercially available grinding control system that includes MPC. With Millstar having the capacity for multivariable control and constraint handling including tank level, hydrocyclone overflow particle size, slurry densities, and hydrocyclone switching for pressure control. Coetzee (2014) presents the behaviour of Mintek's StarCS Robust Non-linear MPC (RNMPC) on a simulated closed circuit grinding system (with hydrocyclones) leading to a flotation feed surge tank. This implementation builds upon RNMPC presented in Coetzee et al. (2010) through practical extensions including state estimation feedback via an unscented Kalman filter (Wan & van der Merwe, 2000), time-delay incorporation in modelling, PV soft constraints, and ability to handle MVs being set to manual (Coetzee, 2014).

Coetzee and Ramonotsi (2016) then demonstrated the incorporation of RTO with the RNMPC controller on an operating platinum processing closed circuit grinding system (with screen classifiers instead of hydrocyclones) leading to a flotation feed surge tank. The RTO performs steady-state optimisation where the objective function can consider complex combinations of targets, constraints, and economic factors. In replacing individual RNMPCs for the mill discharge sump and the surge tank, and through manipulation of their respective water addition and outflows, the combined RNMPC with RTO resulted in improved long-term SP tracking accuracy and stability.

Le Roux et al. (2016) presented MPC for the control of a simulated variable speed grinding mill in closed circuit with a hydrocyclone and incorporating commonly monitored variables. Dynamic inversion was used for the fast-acting mill sump level control and non-linear MPC for slower circuit wide control, with filtering techniques applied for state estimation and noise handling. Both mill fresh feed throughput and hydrocyclone overflow particle size were targeted, with the controller's performance in tracking these individual SPs and rejecting disturbances presented. Botha et al. (2018) then investigated methods to facilitate discrete control actions, as required when bringing individual hydrocyclone in or out of operation. A hybrid non-linear

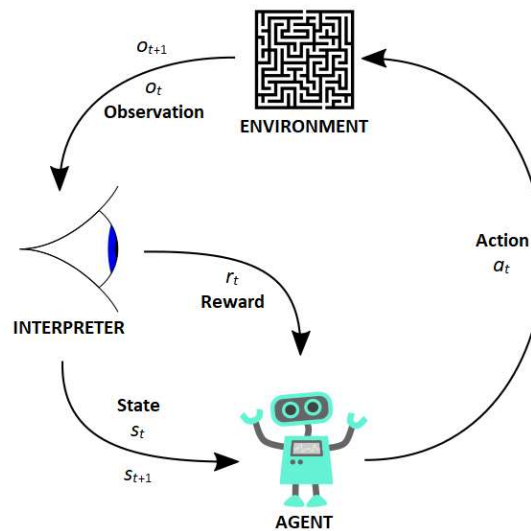
MPC was implemented using genetic algorithms, with the intended benefit of having improved pre-emptive changes to the number of hydrocyclones in operation than if advanced regulatory control was used.

2.6.3 Summary

Currently, hydrocyclone control in the mineral processing industry typically only extends to regulatory control for process stability rather than outright optimisation. Of the advanced control methods, model-based appear to be the most developed, with Mintek producing a commercially available product incorporating MPC for comminution circuit control. The industrial uptake of these model-based methods relies on the successful production of a process dynamics model from available sensor information and its subsequent real-world performance, along with the ever-present factors of cost and willingness to engage with external vendors for proprietary products.

2.7 Reinforcement Learning

RL is another control method applicable to complex systems that has recently found renewed interest following successes with incorporation of deep neural network function approximators. RL is a field of machine learning separate from the likes of supervised learning (learning from labelled data) and unsupervised learning (learning structure in unlabelled data). In RL, the intent is to learn an action taking policy (μ) to maximise a reward signal through interaction with an environment (Sutton & Barto, 2018). RL problems are usually framed as Markov decision processes (MDP). Here an agent (e.g., control system) can interact with, and receive feedback from, the environment (e.g., industrial process). An observation of the environment (o_t) is received (e.g., sensor readings) at each discrete time step (t) and is used to produce a state representation (s_t). Given this, an action (a_t) will then be taken and reward (r_t) received with the subsequent environment observation (o_{t+1}) used to produce the next state representation (s_{t+1}), with a transition probability $p(s_{t+1} | s_t, a_t)$, where the process can then continue from (Lillicrap et al., 2015; Sutton & Barto, 2018). This interaction between agent and environment is illustrated in Figure 2-9.

Figure 2-9*Reinforcement Learning Agent-Environment Interaction*

Note. With representations shown for the current time step (t) and subsequent time step ($t+1$). Adapted from *File:Reinforcement Learning Diagram.svg*, by Megajoice, 2017, Wikimedia Commons (https://commons.wikimedia.org/wiki/File:Reinforcement_learning_diagram.svg).

CC0

1.0.

The reward is that obtained for a single time step, though it is the long-term return over an extended period that the agent should maximise. Value functions are used to quantify the cumulative expected return (J), such as the state-value function $V(s_t)$ being the expected return from the current state or the action-value function $Q(s_t, a_t)$ being the expected return for the action taken from the current state (Sutton & Barto, 2018). The behavioural policies in RL can take the form of a stochastic policy $\mu(a_t|s_t)$ where an action is selected from a probability distribution (given the current state) or a deterministic policy $\mu(s_t)$ where the action is directly selected given the current state (Silver et al., 2014). The process of developing an optimal policy in RL can occur on-policy where the action taking policy is also being optimized, or off-policy where a target policy is being optimized while using data resulting from an alternative behavioural policy (Sutton & Barto, 2018). While RL methods can

incorporate models of the environment's behaviour which can be used for planning a course of action, model-free methods will develop a policy through learning from environment interaction alone (Sutton & Barto, 2018).

2.7.1 Industrial Process Control

Investigations into industrial control with RL have been approached in a variety of ways, progressing with the introduction of new techniques and improved computational power. Hoskins and Himmelblau (1992) investigated using shallow ANNs to form evaluation and action networks for temperature control of a simulated continuous stirred tank reactor, using discrete actions and compared against PID control. Govindhasamy et al. (2005) presented an actor-critic type controller with shallow ANNs for the optimisation of a batch disc grinding process, producing a lower rejection rate than the existing control strategy. Jiang et al. (2018) presented a mineral processing specific RL control design for concentrate and tailings grade control of a simulated single flotation cell. In this case, along with ANN based critic and actor networks, they also incorporate an ANN model to approximate process dynamics from system inputs to outputs. A common thread in these industrial control investigations is the use of the actor-critic structure. The actor learns a policy for controller behaviour at a given time while the critic learns a value function used in critiquing the actor's behaviour, with this interplay facilitating learning (Sutton & Barto, 2018).

In terms of recent comminution circuit related RL control, Conradie and Aldrich (2001) conducted an investigation using a simulated ball mill in closed circuit with a hydrocyclone. A shallow policy neural network was used, with six input variables (including the hydrocyclone overflow particle size SP), a single twelve node hidden layer, and outputs five continuous actions. The reward function was structured to target the hydrocyclone overflow particle size SP along with maximising the mill feed throughput rate, and the policy trained through Moriarty and Miikkulainen's (1996) symbiotic adaptive neuro-evolution genetic algorithm. Conradie and Aldrich (2001) performed trained model testing through observing the response to hydrocyclone overflow particle size SP changes, with unmeasured variability in the form of mill feed particle size and ore hardness trialled. In taking a

circuit wide view, the dominant action learnt by the controller was to manipulate the mill fresh feed throughput to maintain the particle size SP.

During the period in which this research project was undertaken, additional comminution circuit RL control investigations were also published. Hallén et al. (2019) used a simulated milling circuit consisting of a primary mill feeding a secondary mill, and in closed circuit with a classifier. The simulation included several PID loops for control, which could be substituted with RL control to varying degrees depending on the experimental scenario requirements. Proximal policy optimisation (PPO; Schulman et al., 2017) was the RL algorithm used, with three continuous actions and one discrete action (a flap-gate for secondary mill pebble addition from primary mill) available to the control. The reward function was structured to maximise profit (predicted recoverable valuable minerals as a function of particle size and throughput) whilst penalising conditions and actions outside of constraint boundaries. With variable ore properties applied during training and testing, the RL control could demonstrate improvements in attained profit under certain conditions compared to extensively tuned PID control.

Guo et al. (2019) also investigated the use of RL on a ball mill in closed circuit with a hydrocyclone. They used a similar three ANN structure as Jiang et al. (2018), but with deep neural networks for the actor policy and critic along with a recurrent neural network trained on historic operational data to model process dynamics. There are five continuous actions available, being the available SPs for the grinding circuit, and with the actor updated via a clipped policy ratio method as in PPO (Schulman et al., 2017). The reward function was structured to encourage a target hydrocyclone overflow particle size or finer, whilst also maximising the mill feed rate (with a discount factor applied to give a greater priority weighting to the overflow particle size). The reward function also included penalties for any variables exceeding their nominated constraint range.

Guo et al. (2019) performed testing of their trained control policy comparing the calculated daily production performance, in terms of overflow particle size and throughput, against that yielded from a historic operating log. One test method considered what the improvement from the control action would be for one time

step (leading to a simulated subsequent state) against what was performed. The other test allowed successive control actions from simulated subsequent states. Both test methods yielded overall production performance improvements for both metrics, as well as greater proportion of beneficial control actions taken compared to the historic logged actions.

2.7.2 Algorithm Background

Numerous recent model-free continuous action and state space applicable algorithms have also been developed. These include the on-policy asynchronous advantage actor-critic (Mnih et al., 2016) and PPO (Schulman et al., 2017), or the off-policy DDPG (Lillicrap et al., 2015), normalized advantage functions (Gu et al., 2016) and soft actor-critic (Haarnoja et al., 2018). DDPG was selected as the base algorithm for this project, being an actor-critic method that has found success in a range of control tasks including indicated extensions to process control by Spielberg (2017).

Silver et al. (2014) presented deterministic policy gradient (DPG) based algorithms as more efficient alternatives to stochastic policy gradient methods in continuous action space RL. Lillicrap et al.'s (2015) DDPG extends the off-policy actor-critic method of DPG with deep neural network function approximators. For DDPG these are the deterministic policy $\mu(s_t | \mathcal{W}^\mu)$ and action-value function $Q(s_t, a_t | \mathcal{W}^Q)$ which are parametrised with learnable weights (\mathcal{W}). In their work on RL feedback control, Hafner and Riedmiller (2011) presented a related batch learning method neural fitted Q iteration with continuous actions which also use an actor-critic design with neural network function approximators.

Mnih et al.'s (2015) DQN used a number of techniques to facilitate learning with large neural networks, which have been incorporated into DDPG. One technique is a fixed size replay buffer where experienced transitions (s_t, a_t, r_t, s_{t+1}) are stored. A mini-batch of samples (M) are then taken at random from the replay buffer during training, allowing the networks to update on uncorrelated historic transitions rather than recent sequential transitions (Lillicrap et al., 2015; Mnih et al., 2015). DDPG's critic action-value function update occurs using an adaption of Q-learning (Watkins & Dayan, 1992) through minimising the loss shown in Equation 2-7, which considers

the critic network output and the target value (R) determined for each sampled transition (n) in the mini-batch, explicitly (s_n, a_n, r_n, s_{n+1}).

Lillicrap et al.'s (2015) critic update to minimise loss

$$L(w^Q) = \frac{1}{N} \sum_n \left(R_n - Q(s_n, a_n | w^Q) \right)^2 \quad (2-7)$$

In DDPG's case the target value is given by the reward plus the discounted future return shown in Equation 2-8. The discount factor (γ) influences the importance given to future rewards, with $\gamma = 0$ only considering the immediate reward while the closer γ is to 1 the stronger the weighting given to future rewards (Sutton & Barto, 2018). This is also where Lillicrap et al. (2015) apply another technique from DQN (Mnih et al., 2015), where the discounted future return is not based on the currently learnt actor and critic networks, but rather target networks (μ' and Q' respectively). The target networks are copies of the actor and critic networks whose weights update slower than the current actor and critic networks, as in Equation 2-9.

Lillicrap et al.'s (2015) target value for critic update

$$R_n = r_n + \gamma Q'(s_{n+1}, \mu'(s_{n+1} | w^{\mu'}) | w^{Q'}) \quad (2-8)$$

Lillicrap et al.'s (2015) target network soft-update

$$w' \leftarrow \tau w + (1 - \tau) w' \quad (2-9)$$

where τ = soft-update factor (value much less than one)

The use of these soft updating targets stabilise learning of the action-value function (Lillicrap et al., 2015). The actor policy function is then updated to maximise expected return through an application of the chain rule via the action value function

gradient with respect to the actions and the policy gradient with respect to the policy network's weights (Equation 2-10), as proven as the policy gradient in the DPG study (Lillicrap et al., 2015; Silver et al., 2014).

Lillicrap et al.'s (2015) actor update via policy gradient

$$\nabla_{w^\mu} J \approx \frac{1}{N} \sum_n \left(\nabla_a Q(s_n, a | w^Q) |_{a=\mu(s_n)} \nabla_{w^\mu} \mu(s_n | w^\mu) \right) \quad (2-10)$$

DDPG's success in a range of continuous control tasks, included car driving simulation, from simple state descriptions or rendered environment has led to many extensions and adaptations. Examples include Hausknecht and Stone's (2015) adaptation of DDPG for the parameterized action space of robotic soccer, where the requirement is to select a discrete action to perform and its associated continuous parameter(s). Fujimoto, van Hoof, and Meger (2018) extended upon DDPG as the Twin Delayed Deep Deterministic policy gradient (TD3) algorithm to address function approximation errors through various techniques. One such technique was clipped double Q-learning using dual critics, with the lower target network's discounted future return used to train both critics, to counter overestimation bias. Another was the use of a delayed policy and target network update rate compared to the critic, which acts as a way of reducing the influence of the parameterised action-value function variance. Spielberg (2017) extended DDPG for the SP tracking requirements of process control. Training was performed on simulated single-input single-output and 2 x 2 multiple-input multiple-output systems, based on plausible industrial tasks and dynamics. Their trained system's response to SP changes and varying degrees of input/output noise was assessed, with acceptable control yielded.

2.7.3 Summary

RL presents an alternative control method that aims to maximise future reward. Industrial process control investigations have been undertaken, including for complex mineral processing systems, with capabilities progressing with techniques and technology improvements. Lillicrap et al.'s (2015) DDPG algorithm was noted to

possess several appealing qualities for process control. It has an actor-critic structure common to other industrial RL examples and can operate in high-dimensional continuous action and state spaces. It is a model free method, which is of particular interest when a system model may not be feasible to produce, and an off-policy method able to learn from historic transitions.

3 Hydrocyclone Operational State Detection

Section 3.1 and Section 3.2 contains material previously published in “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, pp. 2–9 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier.

3.1 Laboratory Hydrocyclone: Single Source

This proof-of-concept investigation was undertaken to gauge the potential for CNNs to produce a three-state classifier (fan/rope/blocked), using a laboratory hydrocyclone (single source) and a fixed camera.

3.1.1 Data Collection

High-speed video footage of a laboratory hydrocyclone was collected using a FASTCAM Mini UX50 recording at 2,000 FPS, with a 1/4000 s exposure. A Nikon AF NIKKOR 1.8D 50 mm prime lens was mounted on the camera, set at f/2.8 with a distance of 1.64 m between the lens and hydrocyclone apex opening. The camera itself was mounted separately from the closed circuit hydrocyclone system to minimise vibration effects on the video data. The primary light source consisted of two 500 W halogen lights on a tripod to left of the camera, with a secondary light source comprising two 20 W LED lights on a tripod positioned to the right of the camera. A black screen was used to enhance the contrast of the illuminated hydrocyclone underflow. The hydrocyclone cone was 183 mm high with a 10 mm diameter apex opening. The system as detailed above is shown in Figure 3-1.

Figure 3-1*Hydrocyclone Laboratory Experimental Setup*


Note. Located at the Western Australian School of Mines (WASM) Kalgoorlie campus. From “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 3 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Reprinted with permission.

Experimental runs were conducted with quartz slurries consisting of different particle sizes. Runs were executed by progressively increasing the solids concentration in the system, until the hydrocyclone roped. After each incremental increase in the solids concentration, the operation of the hydrocyclone was allowed to stabilise before video data were collected. In addition, the feed valve would be adjusted to modify the hydrocyclone feed flow rate. Each recording period filled the camera’s 16 GB memory, resulting in 19,409 individual frames (approximately 9.7 s worth of footage). The model training and testing frames were then drawn from footage of experimental runs with hydrocyclone feed solids concentration between 5.0%–17.0% by weight (w/w) and inlet gauge pressure 30–80 kPa.

3.1.2 Model Development

Raw image frames from the experimental runs were 12-bit grey scale with a 768 x 768 pixel dimensionality. NumPy (Oliphant, 2006) & Pillow (Clark & contributors, 2010) Python libraries were used to convert the raw images to 8-bit grey scale (i.e., the pixel intensity was rescaled from 0–4,095 to 0–255) and then stacked over three channels to produce an RGB type image of the form expected by the pretrained CNNs to be employed. Moreover, the images were cropped square, approximately centring the hydrocyclone apex in the frame and removing the overflow and support bracket from the original image. Synthetic images to represent a blocked hydrocyclone underflow (i.e., no discharge) were produced by using roping images from another experimental run, but covering the discharge just below the apex with a section of the background from the image.

2,000 consecutive frames were taken from the three experimental runs used to represent the operating states, (i.e., fan, rope and blocked). The first 1,400 frames were used to build a training set, the next 400 frames were used for the validation set, and the final 200 frames were used for the test set. DIGITS (NVIDIA Corporation, 2014) was employed for dataset construction, resizing all images to 256 x 256 pixels. Examples of these frames are shown in Figure 3-2.

Figure 3-2*Example Raw and Training Images*


Note. Raw images (top), and the same images processed for the training dataset (bottom). Adapted from “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, Minerals Engineering, 149, p. 3 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Adapted with permission.

In producing the state detection model, transfer learning was performed using a Caffe (Jia et al., 2014) implementation of VGG-16 (Simonyan & Zisserman, 2014) available on DIGITS (NVIDIA Corporation, 2014), pretrained on the ILSVRC2012 dataset. Mean pixel subtraction was employed, along with cropping to 224 x 224 image size and image mirroring during training. The convolutional layers were unchanged with their pretrained weights frozen, the standard fully connected layers were replaced with a fully connected layer with 15 nodes and ReLU activation, followed by a 50% dropout layer (Srivastava et al., 2014), and finally a three node fully connected layer with softmax activation for classification. Figure 3-3 illustrates

the architecture of the new classification top layers. A training batch size of 42 and stochastic gradient descent was used to train the network, with a base learning rate of 1×10^{-4} and decay factor of 0.1 applied after completion of 33% intervals of the two training epochs. Training and validation accuracies of 100% were achieved, as could be expected for distinctly different conditions.

Figure 3-3

Laboratory Single Source Three-State Classification Network Architecture



Note. Shown from the VGG-16 (Simonyan & Zisserman, 2014) final convolutional block max pooling layer. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

3.1.3 Testing and Discussion

Classification of the 200 test images not seen by the model during training also resulted in 100% accuracy. To interrogate the model further, underflow images from other experimental runs with different feed particle sizes, ore compositions and operating conditions were presented for classification. Of particular interest was a period of unstable operation in which the hydrocyclone underflow was fluctuating between fanning and roping conditions. Further synthetic blockage conditions were also generated from other experimental runs, though with slight modification, such as leaving some of the underflow in the image to give the appearance of splashing slurry.

Figure 3-4 shows the classification results of the trained model on these additional test images. All blockage conditions were correctly identified, even with the presence of noise in the images. The allocation of fanning and roping appears reasonable, considering the model was only trained on a single stable fan period and

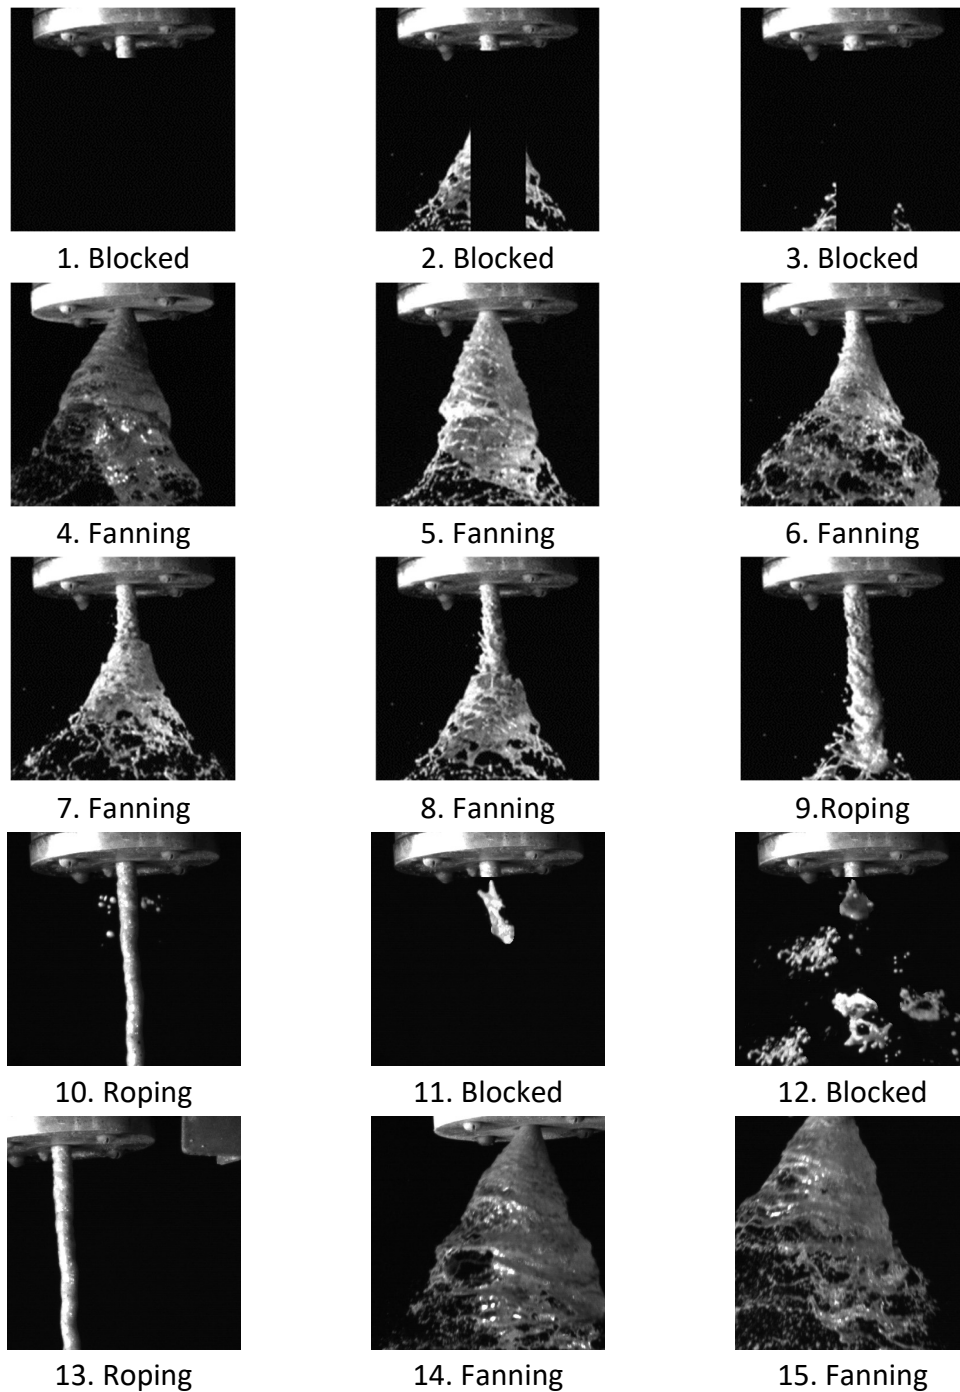
single consistent roping period. The point at which an image begins to be classified as roping could be adjusted simply by training the model with other example images that would also be considered roping.

It should be noted that under these conditions the underflow width at a fixed horizontal plane method proposed by Janse van Vuuren et al. (2011), if extended further to consider blockage conditions and individual state classification, would also have given meaningful results for a number of images. However, they noted particular issues with foreground noise which result in high intensity regions in the image that were difficult to handle through the image processing methods trialled. This could result in erroneous underflow width measurement, and potentially incorrect state detection, if present within the search region (Janse van Vuuren et al., 2011). The assessment of the features within the image as a whole, and their correlation to what constitutes the distinct states, makes CNNs based methods more robust to this type of noise. Figure 3-4 images 10 to 12 present synthetic examples of foreground noise which would have presented a challenge for the fixed horizontal plane method, but have been correctly classified by the CNN based sensor.

Another benefit of the use of CNNs is their ability to handle image translation (Goodfellow et al., 2016). With the method employed by Janse van Vuuren et al. (2011) the sensor is set up to monitor a specific horizontal line of the underflow image, and determine the underflow width within a specific interval search limit of that line (kept narrow to minimise potential for noise to enter and affect that measurement). This means that anything that may shift the section being monitored, such as cleaning or maintenance on the camera, would require confirmation of the sensor's accuracy and adjustment as required. This is not a major issue for a CNN based sensor, as demonstrated by significant image shifts and resulting correct state allocation shown in Figure 3-4 images 13 to 15, thus leading to a more stable long term sensing option when applied in the field.

Figure 3-4

Laboratory Three-State Detector Test Images and Output Classification



Note. Additional test images and their classification returned by the trained model, as indicated by the labels. Adapted from “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 4

(<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Adapted with permission.

Without the need to compare consecutive frames to establish points of motion or smoothing of an underflow width measurement over time with a moving average as in Janse van Vuuren et al. (2011), the CNN based method is able to assess each frame for the probability that it is a member of each class. In this way the unstable transition period prior to continuous roping could be detected if the underflow profile collapses to the extent that it would be detected as roping for just a single frame. Thus, alarm thresholds could be set for the number of roping detections over a given time period to differentiate between the transition period and continuous roping. Figure 3-5 shows an example of an unstable transition period of the laboratory hydrocyclone filmed using a Panasonic Lumix DMC-FT2 hand-held camera at a more typical frame rate of 30 FPS. The achievable analysis frequency will depend on what is the bottleneck of the process such as frame rate or inference processing time, given the available hardware and model employed.

Figure 3-5

Laboratory Hydrocyclone Unstable Transition Period



Note. Three consecutive frames of the laboratory hydrocyclone during an unstable transition period. Adapted from “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 4

(<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier.

Adapted with permission.

This work demonstrates a proof of concept by simply presenting labelled images of hydrocyclone underflows in these three states, CNNs are able to extract meaningful features for their classification. There is thus potential to produce a fixed camera state detector with an ability to generalise to unseen operating conditions and handle a degree of image noise beyond what has been indicated by previous image-based sensors.

3.2 Industrial hydrocyclone: Multiple Sources

This preliminary hydrocyclone investigation was undertaken to gauge the extent to which CNN based sensors can generalise and be robust to more varied industrially relevant conditions. A two-state classifier (fan/rope) was investigated using a variety of industrial and laboratory hydrocyclone footage sources in which the cameras were not mounted to a fixed position.

3.2.1 Data Collection

Video material of fanning and roping operating conditions in hydrocyclones were sourced online and from industry donations. Online sources of hydrocyclone video footage used were CDE Group (2011), 911 Metallurgy Corp. (2016), Cyclone Engineering Jinyang (2016), and Xinhai Mining (2016). With the number of frames used within fair dealing for research and study limits of Australian Copyright Act. Roping footage from the laboratory hydrocyclone system used in Section 3.1, but filmed with a Panasonic Lumix DMC-FT2 hand-held camera, were also collected to expand this class's dataset. The examples available present a range of images of varying quality, different viewing positions, hydrocyclone types and circuit designs, ore types, and the presence of image noise, such as mist and slurry spray. An effort was made to produce a balanced dataset with regard to the number of samples in each class, along with the number of examples from a given source, as summarised in Appendix A. Roping and fanning states only were considered, as an insufficient number of samples related to the blockage of hydrocyclones could be sourced.

3.2.2 Model Development

The video clips were split into individual frames using the OpenCV (Bradski, 2000) Python library (opencv-python package; Heinisuo et al., 2019), rotated so that the apex and discharge were orientated approximately north to south, and cropped square using the NumPy & Pillow Python libraries. Preprocessing of the images were done to ensure that the hydrocyclone apex and discharge were within a reasonable distance from the image edges and that objects such as other hydrocyclones and discharging slurry were largely removed, as at this stage the sensor is being developed to determine the operating state of one hydrocyclone per frame. The resulting dataset contained 322 examples of each class for training and 102 examples of each class for validation. Included in each class's validation set were frames from unseen sources (i.e., two unseen fan sources and one unseen rope source), along with different viewing angles from sources included in the training set, to encourage the resulting classifier's ability to generalise.

With the images resized to 224 x 224 pixels and the network's relevant input preprocessing performed, these images were then passed through the Keras (Chollet & others, 2015a) implementation of ResNet50 (He et al., 2015). This was available via Tensorflow's (Abadi et al., 2015) implementation of Keras and pretrained on the ImageNet dataset. With the classification layer removed this produced a 2,048 value feature vector per image. Scikit-learn's (Pedregosa et al., 2011) C-Support Vector Classification (SVC) was employed on the training dataset with a sigmoid kernel, $C = 4$ error term penalty parameter, and $\gamma =$ scale kernel coefficient resulted in a 93.6% validation set accuracy (13 errors out of 204 images). Indicating that meaningful features were able to be extracted from these industrially sourced images and used in state detection.

In addition, the ability of the CNN itself to classify the images was also evaluated. As previously, the images were resized to 224 x 224 pixels and the network's relevant input preprocessing was performed. In this case Keras' implementation of VGG-19 (Simonyan & Zisserman, 2014), pretrained on the ImageNet dataset and with the fully connected top layers removed, was employed. The weights of this feature extracting part of the network were fixed during training,

and result in 512 7x7 feature maps per image. Image augmentation applied using Keras *ImageDataGenerator* was performed during training to avoid presenting the network with an identical image more than once, effectively expanding the dataset.


Image augmentation involved constrained random position shifts, zooming, rotation, brightness and colour shifts, mirroring and with any blank areas in the image generated filled black as it is irrelevant for discerning the image class. The *ImageDataGenerator* settings are shown in Table 3-1, and examples of augmented image are shown in Figure 3-6.

Table 3-1

Data Augmentation Settings

Variable	Setting
rotation_range	10
width_shift_range	0.1
height_shift_range	0.1
shear_range	0
zoom_range	0.05
brightness_range	0.9, 1.3
channel_shift_range	20
horizontal_flip	True
fill_mode	constant
cval	0

Note. Keras *ImageDataGenerator* (Chollet & others, 2015a) settings used in model development. From “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 9 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Reprinted with permission.

Figure 3-6*Example Augmented Images*


Note. Examples of image augmentation randomly performed by use of Keras *ImageDataGenerator* (Chollet & others, 2015a). Original image top left. Adapted from “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 5 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Adapted with permission.

The classifier portion of the network consisted firstly of applying dropout with a 50% probability, followed by a five node fully connected layer (kept small to limit the model’s capacity) with ReLU activation and L1 and L2 kernel regularisation (with penalties of 0.2 and 0.1 respectively). This regularisation scheme acts to encourage the reduction and zeroing of the layer’s weights, in a similar vein to elastic net (Zou

& Hastie, 2005). The five node fully connected layer was followed by a single node output layer with sigmoid activation, giving an output range 0–1. Figure 3-7 illustrates the architecture of the new classification top layers. Binary cross entropy loss (Equation 3-1) was employed, a training batch size of 46, and with Adam (Kingma & Ba, 2014) optimizer using a learning rate = 0.001 and decay rate = 0.1.

Figure 3-7

Industrial Multiple Source Two-State Classification Network Architecture



Note. Shown from the VGG-19 (Simonyan & Zisserman, 2014) final convolutional block max pooling layer. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

Binary cross entropy loss

$$L(w) = -[y \log \hat{y}(w) + (1 - y) \log(1 - \hat{y}(w))] \quad (3-1)$$

where y is probability true class = 1

(thus $1 - y$ is probability true class = 0)

\hat{y} is the network's output

(Murphy, 2012)

The unconventional approach of applying dropout between the VGG-19 feature extractor and the hidden fully connected layer was employed as a way of adding noise to the classifier input, forcing it not to over-rely on any given features. Owing to the small size of the dataset and similarities between images from the same video footage, data augmentation, dropout, L1 and L2 regularization were all employed to minimise the potential for overfitting and promote the network's ability to generalise.

Random number seeding was employed to limit the run-to-run variability, though with training employing GPU for efficiency thus some variability still occurs (Chollet & others, 2015b). Hyperparameters were selected to give a high validation set accuracy, as an indication of ability to generalise. This model resulted in a 93.6% validation set accuracy (13 errors out of 204 images) after 10 epochs.

3.2.3 Testing and Discussion

For the trained VGG-19 CNN model, of the 13 validation set image misclassifications seven were from 35 frames of industrially donated footage. This system had poor lighting and slurry spray in the footage, however when looking at the specific frames that were misclassified the relatively narrow slurry covered spigot housing and wide flange joining it to the cyclone body were prominent, making a misclassification as roping plausible. By cropping these frames to remove the flange from the image and focusing more on the discharge region, all were then able to be correctly classified as fan state by the model.


A further five of 64 frames from an online sourced roping state footage, which was relatively low resolution, and one of 35 frames from an online sourced fan state footage were also misclassified. These frames did not have any visually obvious reasons for misclassification compared to others from the same footage, though given the improvements yielded by cropping to make the hydrocyclone discharge region more prominent in the frame this was also trialled. This resulted in three of the roping state frames and the single fan state frame being reclassified correctly by the model. Though the model's accuracy could still be refined further, this helps to demonstrate that in practice by initially presenting the best possible images the chance of misclassification even of unseen systems is lowered. This could include camera positioning considerations, cropping to focus on the region of interest, and optimisation of conditions such as lighting as much as practicable.

Since the training and validation datasets consisted of only roping examples for the laboratory hydrocyclone, images of the system operating in fan state were tested to confirm that the classifier could correctly infer this unseen scenario (Figure

3-8). The fan state was correctly classified for both ambient light and illuminated conditions.

Figure 3-8

Industrial Two-State Detector Test Images and Output Classification



Note. Classification of the laboratory hydrocyclone underflow test images by the convolutional neural network (predicted states indicated by the labels). Adapted from “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 5 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Adapted with permission.

As a further test, another laboratory hydrocyclone was filled with water and allowed to drain by gravity. Images of this were passed through the trained network, which classified it as roping. The significance of this being that as roping is an undesirable condition and should ideally be a rare occurrence, it would be difficult to verify if a pretrained image-based sensor could detect the condition, should it occur. Given that this water flow was classified as roping, a simple test could be to run water through the hydrocyclones of new installations in such a way that similar water roping occurs. Images of these flows could then be used to validate the sensor output. If classified as roping, it would be an indication that a slurry rope should be detected correctly. If it is not classified as roping then fine tuning of the classifier may

be required, potentially using these water rope images as a substitute for actual roping footage of the system.

Additional tests on footage donated from the Industry B source, of different hydrocyclones than those seen in the training process, were also undertaken to gauge the model's current performance on these new systems. One set of 886 fan state frames from a non-stationary camera resulted in 213 misclassifications (76.0% accuracy), while another set of 122 fan state frames of a non-stationary camera resulted in 21 misclassifications (82.8% accuracy). This suggests that though efforts have been made to counteract it, overfitting to the training and validation set may still have occurred.

The use of an ensemble method can improve prediction accuracy (reduce generalisation error) through averaging model outputs. This is beneficial as models trained in distinctly different ways are unlikely to make exactly the same errors. One such method would be a majority vote, in which each model assesses a frame and the majority class predicted is the output of the ensemble (Goodfellow et al., 2016). Given that only two distinct models have been developed on the industrial data so far, the average output probability method will instead be used. Here the probability that the image belongs to each class, given by each model, will be averaged to determine the class. Thus, in the event of a disagreement the model which appears more certain in its prediction of the class will determine the ensemble's output, exploiting strengths in interpreting different situations that each model may have developed.

Scikit-learn SVC's ability to employ Platt's scaling to produce a probability estimate from the SVM output (Platt, 2000) was utilised for the previously trained ResNet50 feature extraction model. The roping (class 1) probability for the trained SVM classifier is then averaged with the VGG-19 neural network classifier's activation to give the ensemble's probability and resulting classification. Table 3-2 presents the accuracy results for the three methods on the Industry B test images. Here it appears that the SVM classifier had a significantly better accuracy for the Industry_B_8 footage, and was slightly less accurate for the Industry_B_10 footage. More importantly, when the class probabilities for the two methods are averaged the

resulting accuracy of the ensemble is greater than that achieved by either individual classifier.

Table 3-2

Modelling Method Performance Comparison

Source	Frames	Accuracy (%)		
		CNN	SVM	Ensemble
Industry_B_8	886	76.0	97.9	98.2
Industry_B_10	122	82.8	79.5	84.4

Note. Industry B test frame accuracies for the VGG-19 CNN method, ResNet50 + SVM method, and ensemble method. From “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 6 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Reprinted with permission.

3.3 Industrial Hydrocyclone: Single Source

This section will extend upon the work presented so far in Chapter 3 and Giglia and Aldrich (2020) with a focus on the process and practical considerations of developing an industrial fixed camera hydrocyclone state detector, with three-state (blocked/fan/rope) classification. This would currently be the most widely accessible route for the mineral processing industry, where in-house sensor development could be attempted on collected site specific footage and using open-source models and software. Compared to the three-state detector developed in Section 3.1 and Giglia and Aldrich (2020) on rather visually clean fixed high-speed camera laboratory footage, the industrial footage to be used presents more noise, variability, and state transition periods. Training with fixed camera footage means the model does not have to be as robust to changes in viewing angle, scale, discharge position within frame or the extreme of generalising to a new system compared to the industrial two-state (fan/rope) classifier developed in Section 3.2 and Giglia and Aldrich (2020).

3.3.1 Data Collection

Video footage was donated by an industrial source from an operating mineral processing plant. Usage permissions for the footage does not extend to publication of video frames, thus relevant aspects of frames will be described as well as practical throughout this chapter. A fixed camera was used for each hydrocyclone in two clusters of five hydrocyclones, recording at 20 FPS. The frame dimensions were 3840 x 2160 pixels, and contained the main body, apex, and slurry discharging into launder for each hydrocyclone of interest. A review of the footage found which hydrocyclone had the greatest period of roping and blockage events, this hydrocyclone will be referred to as Cyclone A, and was used for initial model training and testing.

This circuit is typically operated with all hydrocyclones in a cluster either online or offline, and prior to taking the cluster offline it is flushed with water resulting in a clean discharge launder. Thus, footage of the hydrocyclone offline will not be used as a substitute for examples of blocked state, given that the discharge launder is prominent when viewing the hydrocyclone's apex and the clean appearance is not representative of how an underflow blockage would appear in operation. In building the training and validation datasets rather strict conditions were used in selecting and labelling images

- **Fan** state images had a broad fan profile filling a substantial portion of the discharge launder
- **Rope** state images had a consistent central rope-like discharge flow extending from the apex
- **Blocked** state images had little to no material discharging from the apex

In long term operation however, hydrocyclone can exhibit intermediate conditions during periods of instability which can be difficult to classify under these strict labels. Frames from these periods have been included in the test dataset to gauge the trained model's response to these conditions. Examples of these intermediate conditions include

- Periods where the fan briefly collapses to varying degree, but does not form a rope, resulting in a disrupted flow pattern. These were labelled as either fan or blockage depending on the extent of disruption.
- Periods where the broad fan had collapsed, and a rope profile is forming near the apex. These were labelled as rope given the sufficient collapse from the typical fan profile.
- Periods of partial blockage or blockage clearing, often occurring between distinct rope and blocked periods. Here falling material to varying extents and appearance is seen between the apex and discharge launder floor. If there was a largely consistent central rope-like discharge extending a reasonable distance from the apex these were be labelled as rope, otherwise they were labelled as blocked.

External lighting was not applied to the hydrocyclone underflows so varied shadows from nearby structures were present to different extents throughout the day, along with bright/overexposed regions depending on the sun's location. These conditions can result in varying degrees of information loss in the image making the classification task more difficult, with the underflow not visible at all during the night being the extreme example of lost information. For a long-term image-based sensor, applied light to the underflow region, and potentially restricting external lighting sources, would be beneficial in providing a more consistent appearance of the underflow profile. Though the current footage provides the opportunity to assess the degree of information loss due to lighting and shadows that can still be handled by a model trained under these conditions. The summary below briefly describes the lighting appearance of the periods of Cyclone A footage used to build the training (Train), validation (Val), and test (Test) sets, with Appendix B showing the state allocations that make up the datasets.

- **Train A** – Low light in discharge launder from wide-spread shadow
- **Train B** – Strong ambient light with prominent shadow centre / centre-right in discharge launder
- **Val A** – Consistent moderate ambient light
- **Test A** – Consistent moderate ambient light

- **Test B** – Strong ambient light with shadow right side in discharge launder
- **Test C** – Strong ambient light with shadow centred in discharge launder
- **Test D** – Strong ambient light with shadow right side in discharge launder
- **Test E** – Consistent low / moderate ambient light

3.3.2 Model Development


The OpenCV (Bradski, 2000) Python library (opencv-python package; Heinisuo et al., 2019) was used to extract the frames from the video footage, with the Pillow (Clark & contributors, 2010) Python library used to extract 448 x 448 pixels crop regions of interest. For each period of footage to be cropped the coordinates for the centre of the apex tip was recorded, defined as $[x_{\text{apex}}, y_{\text{apex}}]$, with this point indicated in Figure 3-9. As while all footage was from Cyclone A events such as cleaning of the camera lens result in changes in this position within the frame. The degree that the apex tip extends vertically into the frame was chosen (v_{apex}) and horizontal centring given the crop size (i.e., 224 pixels) such that the top left coordinates for the crop position $[x_{\text{crop}}, y_{\text{crop}}]$ is given by $x_{\text{crop}} = x_{\text{apex}} - 224$ and $y_{\text{crop}} = y_{\text{apex}} - v_{\text{apex}}$. For this investigation $v_{\text{apex}} = 120$ pixels was selected as a reasonable distance for the apex to extend into the frame given the crop region size, discharge appearance, and degree of crop position shift to be performed during training and testing. From the fixed crop coordinates $[x_{\text{crop}}, y_{\text{crop}}]$ for each period of footage random integers from nominated ranges could then be added to or subtracted from each coordinate to facilitate the desired degree of crop position shift for an investigation.

The training dataset consisted of Train A and Train B datasets, thus there were 200 training images per class for a total of 600 training images. The validation dataset was made up of only the Val A dataset, thus there were 60 validation images per class for a total of 180 validation images. For each epoch of model training the Train A and Train B images were randomly cropped (with random integers between -40 and 40 pixels added to their fixed crop coordinates) and augmented, resulting in each image being cropped and augmented differently for every epoch. The validation dataset was formed from the Val A images initially being randomly cropped (using the same settings as applied to the training set) and with no image augmentation applied,

resulting in the same unaugmented cropped images being used for validation every epoch.

Figure 3-9

Cyclone A Underflow Discharge Region Illustration



Note. The centre of the apex tip used in crop coordinate calculations is indicated by a red + symbol. The 448 x 448 pixels crop region from the fixed crop coordinates (where $v_{\text{apex}} = 120$ pixels) is illustrated by a red box. The extreme top left coordinates for the random cropping method $[x_{\text{crop}} \pm 40, y_{\text{crop}} \pm 40]$ are illustrated by red r symbols.

The industrial fan/rope sensor in Section 3.2 and Giglia and Aldrich (2020) was intended to be able to generalise to previously unseen hydrocyclone systems, and hence more varied and aggressive image augmentation was applied. For this work the three-state (blocked/fan/rope) sensor was trained and tested against the same hydrocyclone of interest, thus only relevant augmentation to this task was applied. Table 3-3 presents the Keras *ImageDataGenerator* (Chollet & others, 2015a) settings for image augmentation used in training the Section 3.2 and Giglia and Aldrich (2020) industrial sensor and the current design. Given the nature of the fixed camera

mounting point, image rotation and zoom augmentation was not applied. A brightness range of 1.0–1.5 was applied given that one training event was particularly dark and the other had a prominent shadow, thus brightness augmentation was intended to represent condition when illumination was increased from this limit. A channel shift range setting of five was applied to give a more subtle change in the image given that in this instance the sensor is not intended to be applied on different circuits and discharge slurries with vastly different appearances. While a horizontality flipped image would not be encountered by the sensor in practice, this augmentation was still applied to the training dataset to add variability, considering the mirror image still has the discharge appearance of the hydrocyclone operating in the labelled state. Given the variable crop position behaviour previously described, further width and height shifts using *ImageDataGenerator* were not required. Fill_mode and cval were left as the default values as none of the applied augmentations result in shifts outside of the input images' boundaries, and thus filling was not required.

Table 3-3

Data Augmentation Settings Comparison

Variable	Settings	
	Giglia & Aldrich (2020)	Current Design
rotation_range	10	0
width_shift_range	0.1	0
height_shift_range	0.1	0
shear_range	0	0
zoom_range	0.05	0
brightness_range	0.9, 1.3	1.0, 1.5
channel_shift_range	20	5
horizontal_flip	True	True
fill_mode	constant	nearest
cval	0	0

Note. Keras *ImageDataGenerator* (Chollet & others, 2015a) settings used in Giglia and Aldrich (2020) industrial sensor and the current work's model development. Adapted from "Operational State Detection in Hydrocyclones with Convolutional

Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals Engineering*, 149, p. 9 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Adapted with permission.


As in Section 3.2 and Giglia and Aldrich (2020), Tensorflow’s (Abadi et al., 2015) implementation of Keras was used in model development. Also following a similar method, the images (augmented as previously described for the training set) were resized to 224 x 224 pixels and the required input preprocessing performed as expected by the pretrained CNN. With this work’s focus being on applied sensor development the lightweight CNN MobileNet (Howard et al., 2017) was selected. The Keras MobileNet implementation, pretrained on the ImageNet dataset, was used with default network width (alpha = 1.0) and resolution multiplier (depth_multiplier = 1). The output of MobileNet’s final convolutional block (i.e., with global average pooling and final classification related layers removed) was used to extract features, with the weights not updated during training and producing 1,024 7x7 feature maps from each input image. For comparison of the networks’ sizes, the Keras MobileNet implementation with these settings consists of approximately 3.2 million parameters to extract features from the three channel 224 x 224 pixels RGB input image to produce 1024 7x7 feature maps. While the Keras VGG-19 (Simonyan & Zisserman, 2014) implementation, as used in Section 3.2 and Giglia and Aldrich (2020), consists of approximately 20.0 million parameters to extract features from the three channel 224 x 224 pixels RGB input image to produce 512 7x7 feature maps.

The classification network was structured similar to Section 3.2 and Giglia and Aldrich (2020) industrial two-state (fan/rope) detector but modified for this three-state (blocked/fan/rope) detection task, as illustrated in Figure 3-10. Dropout (Srivastava et al., 2014) with 50% probability was applied, then a fifteen node fully connected layer with L1 (0.01) and L2 (0.1) kernel regularisation and ReLU activation. Finally, a three node fully connected layer with softmax activation was used given the three-state classification requirement. Categorical cross entropy loss was used with Adam (Kingma & Ba, 2014) optimizer (learning rate = 1×10^{-4}) and a training batch

size of 30. Training was performed under the conditions described for 10 epochs and resulted in a 97.2% validation set accuracy.

Figure 3-10

Industrial Single Source Three-State Classification Network Architecture



Note. Shown from the MobileNet (Howard et al., 2017) final convolutional block output. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

3.3.3 Testing and Discussion


The images from the Test A to Test E datasets were then evaluated by the trained model, with various cropping methods used as detailed below (each crop method used once per image). Figure 3-11 also shows all the crop position shifts for the random and expanded methods (with respect to the relevant fixed coordinates [x_{crop} , y_{crop}]) that were applied to the test datasets, to aid in visualising the behaviour of the crop methods.

- **Fixed** – All images cropped using the relevant fixed crop coordinates. This can also be thought of as the central crop position of the random cropping method used during training.
- **Random** – All images were cropped by having random integers between -40 and 40 pixels added to their relevant fixed crop coordinates. This being the same cropping method used on the training and validation datasets.
- **Expanded** – All images were cropped by having a random integer between -80 and -41 or 41 and 80 pixels added to one of the relevant fixed crop coordinates, and a random integer between -40 and 40 pixels added to the

other fixed crop coordinate. This resulted in crop positions just outside of the region used during training.

Figure 3-11

Test Images Crop Position Shifts



Note. Crop position shifts with respect to the relevant fixed coordinates [x_{crop} , y_{crop}] for the random and expanded crop methods as applied to the Cyclone A test images.

Three-state confusion matrices for each test dataset and crop position are presented in Appendix C, with Table 3-4, Table 3-5, Table 3-6 showing the overall test three-state confusion matrices for each crop method; fixed, random, and expanded respectively. The initial discussion will be around the classification performance on the Fixed crop method (as this removes a source of variability), before moving on to the behaviour with variable crop positions.

Table 3-4

Cyclone A Fixed Crop Method Combined Testing Three-State Confusion Matrix

Crop: Fixed		Predicted State			
<i>Accuracy 79.8%</i>		blocked	fan	rope	<i>Recall</i>
True State	blocked	868	51	365	67.6%
	fan	19	794	0	97.7%
	rope	22	1	152	86.9%
<i>Precision</i>		95.5%	93.9%	29.4%	

Note. Combined (Test A to E datasets) three-state confusion matrix for the fixed crop method, along with per class precision, recall, and overall accuracy.

Table 3-5

Cyclone A Random Crop Method Combined Testing Three-State Confusion Matrix

Crop: Random		Predicted State			
<i>Accuracy 77.8%</i>		blocked	fan	rope	<i>Recall</i>
True State	blocked	829	45	410	64.6%
	fan	29	784	0	96.4%
	rope	19	2	154	88.0%
<i>Precision</i>		94.5%	94.3%	27.3%	

Note. Combined (Test A to E datasets) three-state confusion matrix for the random crop method, along with per class precision, recall, and overall accuracy.

Table 3-6

Cyclone A Expanded Crop Method Combined Testing Three-State Confusion Matrix

Crop: Expanded <i>Accuracy 67.3%</i>		Predicted State			<i>Recall</i>
		blocked	fan	rope	
True State	blocked	719	65	500	56.0%
	fan	122	668	23	82.2%
	rope	29	3	143	81.7%
<i>Precision</i>		82.6%	90.8%	21.5%	

Note. Combined (Test A to E datasets) three-state confusion matrix for the expanded crop method, along with per class precision, recall, and overall accuracy.

3.3.3.1 Fixed Crop Method Classification

As the fan state is the typical and desired hydrocyclone operating state, the overall test recall of 97.7% for the fan state is a key metric. This presents the correctly predicted fan states as a proportion of the total fan labelled (true state) test examples. This provides an indication of the model's potential for false fault state predictions (blocked and rope both being fault states) and subsequently false alarms. A total of 19 fan labelled test images were misclassified as blocked and none misclassified as rope. The Test D dataset contained 18 of these misclassifications, with these fan images all showing a state of fluctuation (not the uniform broad fan of stable operation) and also not the collapse of discharge into the initial roping appearance. Test B, Test C, and Test E datasets all had no fan state misclassifications, and Test A had a single fan misclassification as blocked without a clear visual cause.

Another key metric is the overall test precision of 93.9% for the fan state. This presents the correctly predicted fan states as a proportion of the total fan state predictions for test examples. This provides an indication of the model's potential for fan state readings to contain misclassified rope or blocked states, with these fault states going undetected. The Test B and Test D datasets had no images misclassified as fan. The Test A dataset had one rope labelled image misclassified as fan, and 10 blocked labelled images misclassified as fan. Of these only one blocked labelled misclassification was from a fluctuating period where in labelling it was deemed to

have more of a blocked appearance than stable fan. The Test C dataset had 10 blocked labelled images misclassified as fan, with two of these being from a fluctuating period.

The Test E dataset had 31 blocked labelled images misclassified as fan, with five of these from a fluctuating period. Within this Test E dataset there were 111 blocked labelled images (with varying degrees of material discharging from the apex) from a period in which spillage from the overflow launder above enters the crop region. This period accounted for 21 of the blocked labelled images misclassified as fan, along with 27 the blocked labelled images misclassified as rope. This period presents a higher proportion of misclassification as fan, though is also a more extreme example of foreground noise than the typical mist and spray found in operation.

In all there are several test images misclassified as fan that are part of unstable fluctuating periods, such as a disrupted profile which can also see fan labelled images classified as blocked by the model as previously discussed, along with a period of atypical/extreme foreground noise. There are still a portion of images that have no visibly obvious potential explanation for the misclassification as fan, thus a degree of missed fault state detection under the current training and test conditions would still occur.

The other aspect to discuss is that of the misclassification of the fault states (blocked/rope). Of the rope labelled states 22 (12.6%) have been misclassified as blocked state. While labelling of rope state was strict in terms of appearance of the discharge this demonstrates a level of disagreement in differentiating these fault states. Of the blocked labelled states 365 (28.4%) have been misclassified as rope state. The majority of these have material to some extent discharging from the apex. This helps to highlight the behaviour of the trained model in assessing these rope to blockage transition states. For this hydrocyclone system apex blockages are not a rare occurrence caused by foreign material obstructing the apex, instead occurring more frequently following conditions resulting in roping. The underflow solids concentration appears to rise to a point where the apex blocks or partially blocks and cycling conditions between roping and blockage occurs.

While firm criteria for roping were used in labelling datasets, the transition states particularly when taken as a stand-alone frame also presented challenges for the labeller. It needs to be acknowledged that trained classifiers may draw a different conclusion on unseen transition states, especially if similar examples did not form part of the training set. This could occur in practice if insufficient examples of the fault states are available, with the model instead trained on footage of water/slurry discharging with a rope-like appearance, as suggested in Section 3.2 and Giglia and Aldrich (2020), and offline hydrocyclone footage in place of blockage example. Of most importance is that a fault state has been identified, even if its true state is debatable, and this is the case with only 4.0% of blocked labelled images and 0.6% of rope labelled images tested misclassified as the non-fault fan state.

3.3.3.2 Variable Crop Method Effect

For exploring the effect of crop position, the focus will be on the overall testing dataset and how the crop method affected the predicted state for the image. As previously noted, each crop method was performed once per test image. Table 3-7 summarises whether the predicted state equals the labelled state, true (T) or false (F), for each test image and crop method (as a percentage of the overall test dataset). For 11.5% of the test images the predicted state did not match the labelled state, regardless of crop method used, suggesting these images as presented were challenging for the trained model. For 58.6% of the test images the predictions matched the labelled state regardless of crop position, thus 70.1% of images were classified consistently (correctly or incorrectly) regardless of crop position. Another 14.1% of the test images had predictions matching the labelled state for the fixed and random cropped images but not the expanded crop. While 3.6% of test images for which the fixed and random cropped images predictions were incorrect, but the expanded crop image was correct. Thus 87.8% of images were classified consistently (correctly or incorrectly) within the crop region incorporated in the training process (as the fixed crop method region is contained within the random crop method region).

Table 3-7*Summary of Crop Method Influence on State Prediction*

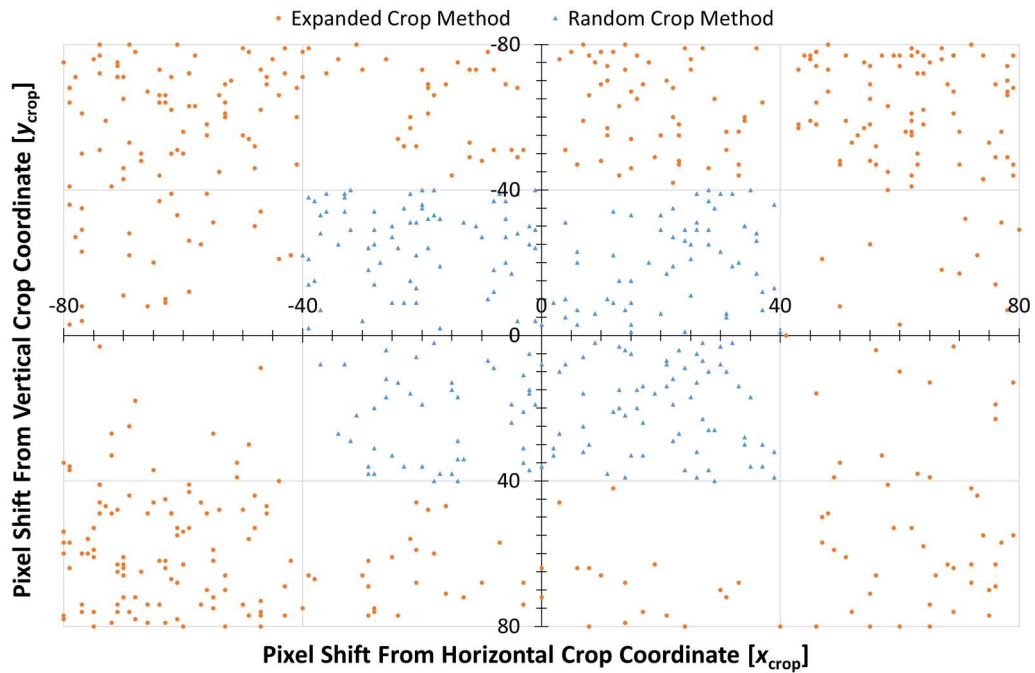
Fixed	Random	Expanded	(%)
T	T	T	58.6
F	F	F	11.5
T	T	F	14.1
F	F	T	3.6
T	F	F	4.1
F	T	T	2.2
T	F	T	3.0
F	T	F	2.9

Note. Shown as a percentage of overall test dataset. Predicted = labelled state (T) and predicted \neq labelled state (F).

Figure 3-12 shows the crop position shifts for the misclassified random and expanded method cropped test images (excluding test images where all three crop methods resulted in misclassification). Figure 3-13 then presents these misclassifications for each 40 x 40 pixels crop region from Figure 3-12, along with misclassification for the fixed crop method, as a percentage of total test images in that region (or the entire test dataset for the fixed crop method case). There appears to be a general increase in misclassifications in the expanded crop region where vertical pixel shift from y_{crop} is less than -40 and greater than or equal to -80, compared to the proportion of misclassifications from the fixed crop method and in the random crop method regions. This vertical shift up in crop position leads to the hydrocyclone being lower down in the cropped frame, thus cutting off part of the region at the bottom of the frame relevant to the underflow discharge appearance. There are also a few expanded crop method regions where horizontal pixel shift from x_{crop} is less than -40 and greater than or equal to -80 which show an increased proportion of misclassifications. This horizontal shift left in crop position leads to more of the launder lip (often covered in solids build-up) being brought into the frame and cutting off some of the discharge launder.

Figure 3-12

Misclassified Test Images Crop Position Shifts



Note. Crop position shifts, with respect to the relevant fixed coordinates $[x_{crop}, y_{crop}]$, for the random and expanded crop methods' test images which were misclassified (excluding test images where all three crop methods resulted in misclassification).

Figure 3-13

Region Percentage of Misclassified Crop Position Shift Test Images

27%	24%	36%	30%	Expanded Crop Method Random Crop Method Fixed Crop Method 9%
21%	13%	11%	10%	
13%	8%	12%	10%	
31%	16%	12%	11%	

Note. Percentage of misclassifications for the respective 40 x 40 pixels regions presented in Figure 3-12. Misclassification count is as used to produce Figure 3-12, representing misclassification associated with the crop position shift and thus excluding test images where all three crop methods resulted in misclassification.

Misclassification percentage is then based on the total test images in that 40 x 40 pixels region (as shown in Figure 3-11). Also shown is the percentage of misclassifications for the fixed crop method (excludes test images where all three crop methods resulted in misclassification) over the entire test set.

Incorporating crop position shift in the training process appears worthwhile to increase the variability of images presented and give a level of robustness to movement within a known region. Monitoring and restricting the applied sensors crop position within the region explored during training would be sensible. Though in this case a 40 pixels shift is less than 10% of the cropping frame's 448 pixels width and height, the introduction of features or appearance changes not experienced during training can lead to a higher rate of misclassification. Ideally, holding the crop position as a fixed region with respect to the centre of the hydrocyclone apex tip would be desirable as it removes a source of variability. The central point of the crop position shift range used in training would be a reasonable point to select, allowing the largest potential shift in any direction to occur and still be contained within the region used during training. Though misclassifications can still occur, as the 2.2% of misclassified test images occurring with the fixed crop method but not the random or expanded crop methods shown in Table 3-7 demonstrate, the consistency a fixed crop region provides is still desirable.

3.3.3.3 Fault State Detection

Considering the previous discussion of the issues encountered by the trained model to discriminate between rope and blocked states, particularly during transition periods which were also challenging for a human labeller, the fault detection performance of the model on the test dataset for each crop method will be examined. Table 3-8, Table 3-9, and Table 3-10 reduces the respective three-state confusion matrices of Table 3-4, Table 3-5, and Table 3-6 to two-state confusion matrices fan versus fault (fault including blocked and rope states). The Matthews correlation coefficient (MCC; Matthews, 1975) was also calculated for each crop method for the now binary classification task (Equation 3-2), as it considers the

proportion of classes and their prediction outcomes to produce a more robust performance score than overall accuracy (Chicco, 2017).

Matthews correlation coefficient (Matthews, 1975) calculated from the confusion matrix values as in Chicco (2017)

$$\text{MCC} = \frac{T_{\text{fan}}T_{\text{fault}} - F_{\text{fan}}F_{\text{fault}}}{\sqrt{(T_{\text{fan}} + F_{\text{fan}})(T_{\text{fan}} + F_{\text{fault}})(T_{\text{fault}} + F_{\text{fan}})(T_{\text{fault}} + F_{\text{fault}})}} \quad (3-2)$$

where T_{fan} = correct fan predictions
 F_{fan} = incorrect fan predictions
 T_{fault} = correct fault predictions
 F_{fault} = incorrect fault predictions

With the model's outputs simplified to the process critical hydrocyclone operational state fault detection, the current performance of the model within the crop region experienced during training appears strong. For the fixed and random crop method tests the performance metrics are comparable, with only a slightly improved overall performance indicated when the fixed crop position is used, as best illustrated by the slightly higher MCC of 0.933 against 0.928. The expanded crop method tests indicate a 90.6% accuracy and MCC of 0.794, further suggesting that there is distinct drop in performance for crop regions outside of that used in training. Of note is the 82.2% fan state recall, thus 17.8% of fan state labelled images were incorrectly classified as fault states and would lead to higher false alarm rates.

Table 3-8

Cyclone A Fixed Crop Method Combined Testing Two-State Confusion Matrix

Crop: Fixed		Predicted State		
<i>Accuracy 96.9%</i>		fan	fault	<i>Recall</i>
True State	fan	794	19	97.7%
	fault	52	1407	96.4%
<i>Precision</i>		93.9%	98.7%	MCC 0.933

Note. Combined (Test A to Test E datasets) two-state confusion matrix for the fixed crop method; along with per class precision and recall, overall accuracy, and Matthews correlation coefficient (MCC).

Table 3-9

Cyclone A Random Crop Method Combined Testing Two-State Confusion Matrix

Crop: Random		Predicted State		
<i>Accuracy 96.7%</i>		fan	fault	<i>Recall</i>
True State	fan	784	29	96.4%
	fault	47	1412	96.8%
<i>Precision</i>		94.3%	98.0%	MCC 0.928

Note. Combined (Test A to E datasets) two-state confusion matrix for the random crop method; along with per class precision and recall, overall accuracy, and Matthews correlation coefficient (MCC).

Table 3-10

Cyclone A Expanded Crop Method Combined Testing Two-State Confusion Matrix

Crop: Expanded <i>Accuracy 90.6%</i>		Predicted State		<i>Recall</i>
		fan	fault	
True State	fan	668	145	82.2%
	fault	68	1391	95.3%
<i>Precision</i>		90.8%	90.6%	MCC 0.794

Note. Combined (Test A to E datasets) two-state confusion matrix for the expanded crop method, along with per class precision and recall, overall accuracy, and Matthews correlation coefficient (MCC).

3.3.3.4 Applied Sensor Considerations

During initial setup of an industrial image-based hydrocyclone state detector consideration to reduce the variability in the footage over time would be beneficial. Applied lighting to the hydrocyclone underflow region would aid in minimising illumination variability, particularly shadows resulting in localised information loss, along with improving contrast with the background. Selection of the camera position and angle so that the resulting crop region contains less visual variability from other sources would also be beneficial. For example, during laboratory hydrocyclone state detector work in Section 3.1 and Giglia and Aldrich (2020), the camera was positioned perpendicular to the vertical hydrocyclone resulting in a crop region containing the hydrocyclone discharge but no flowing and splashing slurry from within the discharge tank.

Additionally, “a clear view of the hydrocyclone’s discharge would generally be required. Spigot socks or housings that obscure a view of the fan profile may make it more difficult to determine the differences between the states” (Giglia & Aldrich, 2020, p. 6). This lack of a reliable external indication of operating state could also be extended to systems where the spray angle leads to a particularly narrow fan profile. In this case, it becomes more challenging to visually discern from the rope state and would likely result in more frequent misclassifications. If typical operation sees a

narrow spray profile than other sensing methods not reliant on the underflow characteristics may be more suitable, such as electrical tomography based (Section 2.3.1) or acoustic based (Section 2.3.2) sensors.

The model trained under the current conditions has not achieved perfect classification accuracy, even for fault detection alone. It was worth considering how the model's output will could be used for alarm thresholds and to form the operating state output that is visible and used by the operators and control system. As the model assesses each frame to output the most probable operating state, one option would be to assess the outputs over a fixed period. For example, if the outputs produced each second were to be assessed, then in this case there would be up to 20 individual frames, if the available hardware allows for all the 20 FPS of video footage to be analysed. A threshold could then be set for the number of frames or consecutive frames in this period to produce a fault alarm (both blocked and rope counted to better handle for transition periods), along with the number of frames or consecutive frames to signify the operating state output for the time period. The number of frames for these thresholds would be set based on balancing the risk of false alarms with missed fault detection. Even with this frequency of model output assessment, the operating state output reading update rate would still be in the order of other commercially available operational state detection methods discussed in Section 2.3.

It is also worth considering the effect of analysing the underflow footage at a lower rate. This would reduce the probability of detecting, and thus responding to, the brief state changes (intermediate periods) that often occur before operating conditions lead to an extended state change. Also with a lower analysis rate, and thus the more time represented by the image analysed, there is more weight being given to that output as there is less data available to compare against to counteract a misclassification. This all factors into the decision on what time period should be used for the alarm thresholds and operating state output, given the desired responsiveness.

If in practice, it is found that consistent features are present at times that result in extended periods of misclassification then another option to improve the

classification accuracy would be through using a model ensemble method. Through the development of additional classification models, such as training with a different pretrained CNN as a feature extractor, each models' probability outputs could be averaged or a majority vote of proposed class could be used to produce the ultimate operational state output for a frame, as discussed in Section 3.2.3 and Giglia and Aldrich (2020). While this method can produce accuracy improvements compared to the individual classifiers, it requires additional computational resources and thus could lower the overall frame rate able to be processed (depending on available hardware).

Both the assessment of frames over time and the use of ensemble methods for individual frame classification can be used in conjunction to try and produce the most reliable state detection classifier under all conditions. It would be worthwhile to have testing footage available to assess the performance of a trained classifier, to give an indication of relevant thresholds and if training additional classifiers may be necessary. At a minimum additional fan operating footage, which should be readily available, should be used to test fan state recall. If no true fault state footage is available then offline hydrocyclone footage as a substitute for blockage, along with artificial water or slurry rope footage, as suggested in Section 3.2.3 and Giglia and Aldrich (2020), would also be worth trialling.

3.3.3.5 Generalisation

The training and validation of the state classifier model was only performed on a small dataset from a single hydrocyclone (Cyclone A) and under visually challenging conditions (particularly due to the lack of dedicated lighting). This was undertaken to gauge the performance potential for a classifier produced under these conditions, though it is also of interest to see if a classifier trained in this way could generalise to other hydrocyclones in the circuit. A period of footage from a hydrocyclone which appeared most similar to Cyclone A in terms of viewing angle and launder position was used to produce test dataset Cyclone B. Frames were cropped using the Fixed crop method, with the resulting appearance having a similar launder lip position in the frame as Cyclone A though with very prominent shadowing covering large portions of the frame. Table 3-11 and Table 3-12 shows the Cyclone B

fixed crop test dataset results in terms of the three-state (blocked/fan/rope) and two-state (fan/fault) confusion matrices, respectively. The fault detection performance on this small test dataset is strong, with no misclassifications as fan, however the fan state recall is only 67.4% with all misclassifications as blockages.

Table 3-11

Cyclone B Fixed Crop Method Testing Three-State Confusion Matrix

Crop: Fixed		Predicted State			
<i>Accuracy 72.1%</i>		blocked	fan	rope	<i>Recall</i>
True State	blocked	39	0	6	86.7%
	fan	45	93	0	67.4%
	rope	9	0	23	71.9%
<i>Precision</i>		41.9%	100%	79.3%	

Note. Cyclone B test dataset three-state confusion matrix for the fixed crop method; along with per class precision, recall, and overall accuracy.

Table 3-12

Cyclone B Fixed Crop Method Testing Two-State Confusion Matrix

Crop: Fixed		Predicted State		
<i>Accuracy 79.1%</i>		fan	fault	<i>Recall</i>
True State	fan	93	45	67.4%
	fault	0	77	100%
<i>Precision</i>		100%	63.1%	MCC 0.652

Note. Cyclone B test dataset two-state confusion matrix for the fixed crop method; along with per class precision and recall, overall accuracy, and Matthews correlation coefficient (MCC).

Another period of footage from a hydrocyclone with a different viewing angle and launder position was used to produce test dataset Cyclone C. Frames were

cropped using the fixed crop method, with the resulting appearance having the launder lip on the right side of the frame (Cyclone A's launder lip is on the left) and prominent shadowing on the centre and lower left of the frame. Table 3-13 shows the Cyclone C fixed crop test dataset results in terms of the three-state (blocked/fan/rope) confusion matrix. A critical issue with the Cyclone C test footage results is that no images were predicted as fan state so in this case the trained classifier has failed to generalise to this new footage. While the extent of shadow present in both these test datasets may go some way to explaining the lower test performance compared to Cyclone A, the poor fan state detection on the Cyclone C test dataset compared to Cyclone B shows that the ability to generalise when trained on a single hydrocyclone is limited. With consistent lighting and fixed camera positions selected to result in a more similar appearance for frames presented to the classifier, it would be of interest to investigate whether training individual classifiers per hydrocyclone or producing a single classifier from the combined dataset would result in improved final performance. This would be of particular interest if the larger combined dataset opens up the option of fine-tuning convolutional layers.

Table 3-13

Cyclone C Fixed Crop Method Testing Three-State Confusion Matrix

Crop: Fixed		Predicted State			
<i>Accuracy 32.2%</i>		blocked	fan	rope	<i>Recall</i>
True State	blocked	56	0	20	73.7%
	fan	38	0	102	0%
	rope	2	0	21	91.3%
<i>Precision</i>		58.3%	-	14.7%	

Note. Cyclone C test dataset three-state confusion matrix for the fixed crop method; along with per class precision, recall, and overall accuracy.

3.3.4 Image-Based Sensor Health Check

While the cameras monitoring the hydrocyclones have fixed mounting points, their brackets allow for adjustment and thus the position of the hydrocyclone within the frame was found to change over the extended periods of footage provided. The manual cleaning of the camera lens was noted as being an event that can result in significant movement of the camera. The state classifier training was performed using variable crop locations containing the hydrocyclone apex and discharge, and thus should have a degree of robustness to movement. Though as an applied sensor would be sampling from a set crop position it would be worth having a health check that the apex and discharge are expected to be acceptably contained within the set crop region.

The more challenging task of object detection could form part of a general moving camera state detector, where confirming the presence and location of a hydrocyclone's apex may be required in producing a reliable operational state classification. Considering in this case the camera mounting point and hydrocyclone structure are fixed, the simpler task of tracking a reference point/region over time would be sufficient. Here tracking is not used to follow a moving object within the frames, but rather to see how the field of view has moved in relation to the object. With the key physical structures and camera fixed, changes in the reference target's scale, rotation and deformation would not likely occur. Changes in the target's appearance due to lighting variation (especially shadows), mist/spray from the operating hydrocyclones, slow build-up of solids on surfaces, and the transition of the camera from colour day mode to a green hued infrared night mode would all be challenges for the tracking algorithm to handle.

3.3.4.1 Tracker Selection

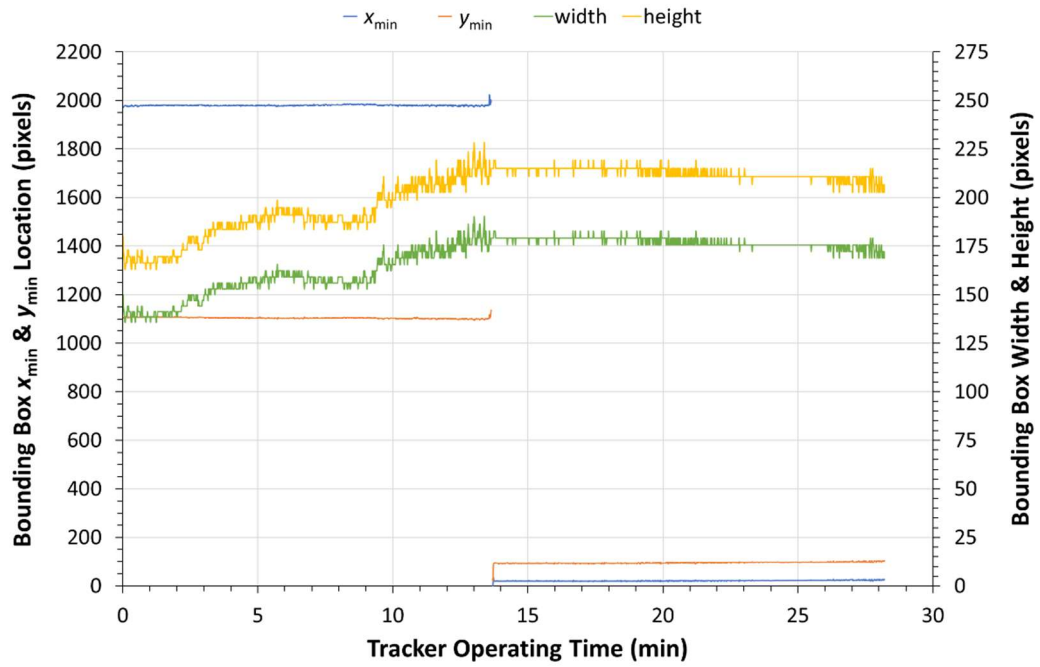
Two OpenCV (Bradski, 2000) tracker implementations, available on the `opencv-contrib-python` package (Heinisuo et al., 2020), were trialled on a period of Cyclone A footage known to contain camera lens cleaning that resulted in a change in field of view. Tracker 1 was *TrackerCSRT* being an implementation of discriminative correlation filter with channel and spatial reliability (Lukežič et al., 2018), and Tracker 2 was *TrackerMOSSE* being an implementation of minimum output sum of squared

error (MOSSE) correlation filter (Bolme et al., 2010). The trackers were initialised with a target region of interest on the first frame of the video footage, this region extended from the tip of the hydrocyclone's apex to part way up the lower cone (including a bracket on the lower cone). The target region is defined by a bounding box (1970, 1105, 150, 180), consisting of the top left corner coordinates x_{min} and y_{min} along with the box's width and height (all expressed in pixels). The tracker was then updated on every 20th frame. Considering the video frame rate is 20 FPS, this monitoring frequency should be more than sufficient to capture movement of the target region within the frame. The OpenCV tracker update for a new frame returns true if the target has been found, along with the proposed bounding box descriptors, otherwise false is returned with all bounding box descriptors zero.

The resulting bounding box descriptors over time for the footage with lens cleaning and subsequent field of view shift for Tracker 1 is shown in Figure 3-14 and for Tracker 2 is shown in Figure 3-15. At approximately 13.5 min into the footage cleaning of the camera lens begins, with both trackers responding with changes in the bounding box x_{min} and y_{min} coordinates as the target region moves within the frame. During cleaning, the camera lens becomes fully obscured and the target is lost, with both trackers returning false and a break shown in the bounding box descriptor plots. Tracker 2 recovers from the obscuration when the target is again visible and continues tracking it. Tracker 1 does not successfully recover from the obscuration and subsequent bounding box descriptors no longer represent the intended target. Prior to the tracking failure, the bounding box descriptors for Tracker 1 are also more variable than Tracker 2, though the scale and structure of the intended target region is not changing during this period.

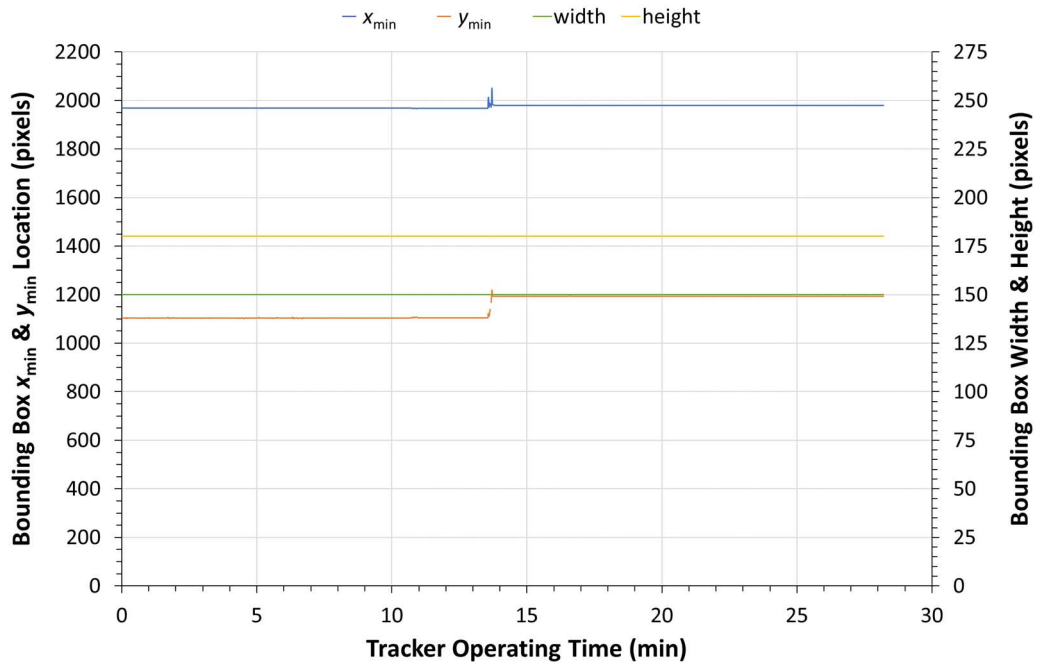
Figure 3-14

Tracker 1 Bounding Box Descriptors During Lens Cleaning



Note. OpenCV *TrackerCSRT* (Heinisuo et al., 2020) bounding box descriptors over time for footage with lens cleaning and subsequent field of view shift.

Figure 3-15

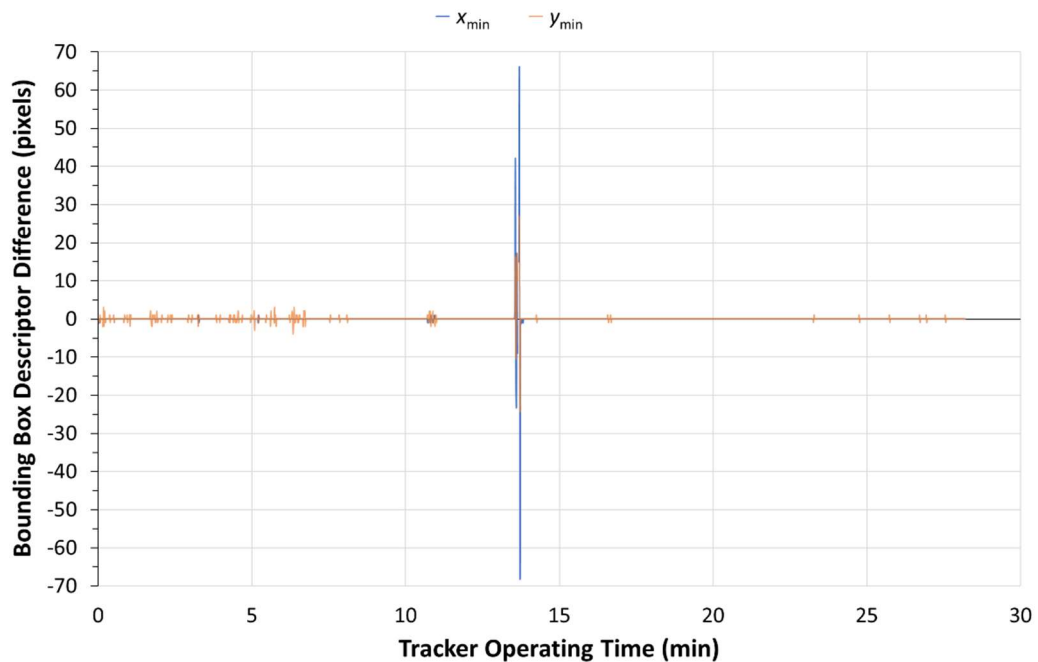
Tracker 2 Bounding Box Descriptors During Lens Cleaning

Note. OpenCV *TrackerMOSSE* (Heinisuo et al., 2020) bounding box descriptors over time for footage with lens cleaning and subsequent field of view shift.

Given the successful recovery from obscuration and more stable bounding box descriptors of Tracker 2, this implementation of *TrackerMOSSE* was investigated further for its potential use in an image-based sensor health check application. With the bounding box width and height remaining constant throughout the tracking process, Figure 3-16 shows only the change in bounding box coordinates between consecutive frames assessed. Minor shifts (in this case within ± 4 pixels) occur during tracking, resulting from causes such as movement/vibration of the system and potentially drift of the tracking region. A significantly larger shift in bounding box coordinates than otherwise observed then occurs during camera lens cleaning in this footage.

Figure 3-16

Tracker 2 Bounding Box Descriptor Difference During Lens Cleaning



Note. OpenCV *TrackerMOSSE* (Heinisuo et al., 2020) bounding box location x_{min} & y_{min} descriptor difference (between current and previously assessed frame) over time for footage with lens cleaning and subsequent field of view shift.

3.3.4.2 Long-Term Stability


To assess the longer term *TrackerMOSSE* behaviour on this system approximately 22.5 hr of continuous Cyclone A footage was monitored, again initialised on the region (1970, 1105, 150, 180) and with the tracker updated on every 20th frame. This was performed to assess response to different visual challenges as previously discussed, along with behaviour during different operating conditions and hydrocyclone start-up and shutdown. Again, with the bounding box width and height remaining constant throughout the tracking process, Figure 3-17. shows the change in bounding box coordinates between consecutive frames assessed. Key events noted are as follows

- a. y_{min} shifts due to hydrocyclone jumping/vibrating on start-up

- b. Two periods of tracking failure returned during lens cleaning. This cleaning instance did not result in a field of view shift
- c. A single frame shift with x_{\min} difference = 9 pixels and y_{\min} difference = 42 pixels, before returning to previous values. No clear cause for the single frame shift, though there was mist/spray occurring in the target region during that period
- d. A single frame tracking failure which resulted in no position shift when the target was subsequently located, and no clear visible cause for the failure
- e. A single frame shift with x_{\min} difference = -1 pixels and y_{\min} difference = 4 pixels, before returning to the previous values, with no clear cause for a single frame shift
- f. x_{\min} and y_{\min} shift due to hydrocyclone jumping/vibrating on start-up
- g. x_{\min} and y_{\min} shift due to hydrocyclone jumping/vibrating on start-up

Figure 3-17

Long-Term Bounding Box Descriptor Difference



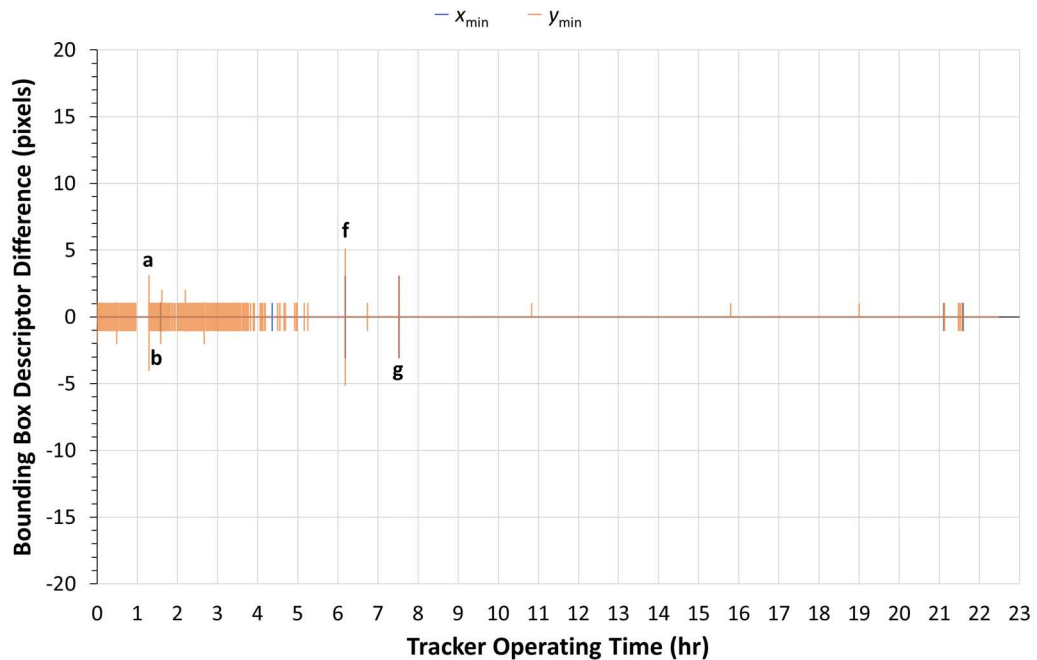
Note. OpenCV *TrackerMOSSE* (Heinisuo et al., 2020) bounding box location x_{\min} & y_{\min} descriptor difference (between current and previously assessed frame) over time.

Extended tracking performance analysis with key events marked. Initial tracking region (1970, 1105, 150, 180).

While most notable shifts and tracking failures were reasonable and explainable, the single frame shifts and failure were more challenging. As the hydrocyclone's apex (which was included in the target region) is within the discharge launder and thus more affected by lighting changes and spray/mist caused by the discharge, a revised target region was considered. The target region was adjusted to mainly include the bracket on the lower cone (1930, 1020, 225, 150), which was further out of the discharge launder and more consistently illuminated. The tracking process was then repeated with the revised target region, with Figure 3-18 showing the change in bounding box coordinates between consecutive frames assessed. The start-up and lens cleaning related points of interest were still present when initialised on this new target region, while the tracking was more stable with none of the unexplained single frame shifts and failure occurring. This demonstrates improved tracker reliability by using a reference point expected to have a more consistent appearance over the long term. A potential issue may be that if a localised obscuration is occurring in the apex and discharge crop region used by the state detector this is less likely to be identified (through the tracker returning a false signal) if the reference point is further away from this region.

Figure 3-18

Long-Term Bounding Box Descriptor Difference With Revised Tracking Region



Note. OpenCV *TrackerMOSSE* (Heinisuo et al., 2020) bounding box location x_{min} & y_{min} descriptor difference (between current and previously assessed frame) over time. Extended tracking performance analysis with key events marked. Revised initial tracking region (1930, 1020, 225, 150).

The change in the bounding box descriptors over the extend tracking period for the two target regions trialled are shown below.

- Initial (1970, 1105, 150, 180), Final (1970, 1110, 150, 180)
- Initial (1930, 1020, 225, 150), Final (1931, 1023, 225, 150)

These changes suggest that there may be drift in the tracker over time, which may require monitoring over an extended period and correction when required. Bolme et al.'s (2010) proposal to counteract drift was to occasionally recentre the MOSSE tracker's filter with the initialisation frame, if the target's appearance is relatively stable. Thus, if drift is found to be excessive this could potentially be incorporated as an automated method, provided there is little variability in the

region selected to be tracked. Matching the recentring schedule to the timing of the initialisation frame (i.e., same time of day), especially if the only lighting source is natural, and selecting a target region that is visibly consistent regardless of operating state would be some of the considerations in trialling automated recentring.

3.3.4.3 Applied Sensor Considerations

From this analysis the use of an appropriate tracking algorithm, such as OpenCV's *TrackerMOSSE* (Heinisuo et al., 2020) implementation, would be useful as a health check system associated with a fixed camera state detector or other image-based sensor. With an appropriate target region selected considering proximity to the state detector crop region (for potential in detecting local obscuration) and long-term appearance stability (for improved tracking reliability) alarm structure and potential automated fault responses could be considered such as

- Log alarm for frames returning false for target detected.
- Log alarm for bounding box coordinate differences to previous frame greater than a certain threshold, dependant on the system and desired sensitivity. For the system assessed, coordinate differences beyond ± 5 pixels could be logged so that operating and start-up movement/vibration are ignored.
- Higher level excessive movement alarm if a given number of frames in a fixed time period return bounding box coordinate differences outside of a set threshold.
- Higher level or critical tracking failure (possible obscuration) alarm if a given number of frames in a fixed time period return false for target detected.
- Higher level alarm if change in average of individual bounding box coordinates over two close time periods is above a given threshold, to capture excessive shift in the field of view. Critical alarm if change expected to have resulted in hydrocyclone apex and discharge not within a valid region of the state detector crop area.
 - Given the fixed camera mounting point and hydrocyclone structure (constant scale), and consistent bounding box dimensions with this implementation of *TrackerMOSSE*, the shift in the bounding box top left corner coordinates could be used to automatically adjust the

state detector crop position following cleaning or camera maintenance (if deemed sufficiently reliable).

- Higher level alarm if change in average of individual bounding box coordinates over widely spaced time periods has occurred, and not already been compensated for. This would be to attempt to detect long-term tracker drift.
 - Automatically recentring of the tracking region based on the initialisation frame could be performed (if deemed sufficiently reliable).

3.4 Summary

CNNs were investigated for their potential to extract meaningful features from hydrocyclone underflow video footage for the classification of operational states. Initially, a three-state (blocked/fan/rope) hydrocyclone operation state detector was trained on fixed camera footage from a laboratory hydrocyclone, using a pretrained version of VGG-16 (Simonyan & Zisserman, 2014) as a feature extractor. The resulting model demonstrated a high classification accuracy and indicated robustness to image noise and camera movement.

A two-state (fan/rope) hydrocyclone operation state detection was then investigated, given the lack of available blocked state examples, and trained on non-fixed camera footage from multiple sources (industrial and laboratory). Misclassification on the validation set of a neural network based model, with a pretrained version of VGG-16 (Simonyan & Zisserman, 2014) as a feature extractor, were found to reduce when the crop position was adjusted to focus more specifically on the underflow regions. Testing on images of water draining from a laboratory hydrocyclone were classified as roping, with such artificial fault state examples potentially useful in training or testing newly installed sensors. Subsequent testing on new fan state footage produced a lower accuracy than that of the validation set used during training. The incorporation of a secondary model, produced using a pretrained version of ResNet50 (He et al., 2015) as a feature extractor followed by a SVM classifier, into an ensemble model resulted in improved testing set accuracy compared to either individual model.

A three-state (blocked/fan/rope) hydrocyclone operation state detector was trained on fixed camera footage from an operating mineral processing plant. This footage was particularly challenging as no additional lighting other than ambient light was applied to the hydrocyclone discharge region, thus presenting a wide extent of visual variability especially due to the presence of shadows. A small dataset was also used for training and validation, with a similar feature extraction and ANN classifier structure as used previously in this work and in Giglia and Aldrich (2020). Though a pretrained version of the efficient MobileNet (Howard et al., 2017) model was selected as the feature extractor. Under these conditions the trained classifier presented strong fault (blocked and rope) versus non-fault (fan) state discernment, particularly within the variable crop region encountered during training. Discernment between the fault states was not as strong, particularly as the test dataset was made up of many transition states between rope and blocked.

The ability for the model trained on a single hydrocyclone to generalise to others within the circuit was limited. Practical considerations for fixed image-based sensor development around minimising variability (through consistent applied lighting and monitoring camera movements) along with model ensemble and temporal information handling for improved accuracy and reliability were also addressed. Of note OpenCV's *TrackerMOSSE* (Heinisuo et al., 2020) implementation was found to be useful in a fixed camera image-based sensor health check application for monitoring camera field of view movement and obscuration.

4 Hydrocyclone Underflow Particle Size Inference

4.1 Laboratory Hydrocyclone

This proof-of-concept investigation was undertaken to gauge the potential for CNNs to infer underflow particle size information from video footage of a laboratory hydrocyclone.

4.1.1 Data Collection

The experimental setup (Figure 4-1) and filming process was as detailed in Section 3.1.1 and Giglia and Aldrich (2020) for high-speed video recording of a laboratory hydrocyclone's underflow. The selection of the f-number for the lens

(f/2.8) was of particular relevance to the following investigation as it aimed to strike a balance between sharpness, depth of field and lighting requirements for high-speed filming of the underflow discharge fan. The frame rate (2000 FPS) and exposure (1/4000 s) were also selected considering the available lighting and whilst attempting to minimise motion blur.

Figure 4-1

Hydrocyclone Laboratory Experimental Setup Viewed From Camera Perspective



Note. Shown from the high-speed camera's view position. Located at the WASM Kalgoorlie campus.

Crushing and milling of a quartz ore source was performed to yield the two different feed particle size distributions used in the experimental runs. Initially 30 kg of water was added to the hopper, and then whilst running on water the feed valve position was adjusted so that the gauge pressure was 100 kPa. The feed material was then added and after a period of stabilisation (at least 2 min) high-speed video footage of the underflow was taken, followed by sampling of the underflow. Sampling was performed using a large beaker to cut the entire underflow stream whilst minimising sample losses through slurry splashing. The feed valve was then adjusted to reduce the hydrocyclone feed flow rate and pressure, then again following a period of stabilisation the discharge was filmed and sampled. The feed


flow rate was then increased, additional solids added, and the above procedure repeated. Appendix D contains specific details of the experimental runs performed. The lights were cleaned as required during this process to remove build up caused by slurry splash and mist. The underflow samples were oven dried for particle size analysis via dry sieving, which was performed as soon as practical to minimize the risk of sample loss. No external assistance was sought, with the entire experimental procedure undertaken by one person to minimise sources of variability and error.

The use of a video camera in this experimental design addressed some of the limitations and recommendations made in Uahengo's (2014) still image-based hydrocyclone investigation. The underflow sampling can be performed whilst the camera is recording, thus image capture and sampling is occurring simultaneously. Also, the video footage allows for sequential images to be incorporated into the investigation.

4.1.2 Footage Classification

4.1.2.1 Model Development

As detailed in Section 3.1.2 and Giglia and Aldrich (2020), NumPy (Oliphant, 2006) and Pillow (Clark & contributors, 2010) Python libraries were used to convert the individual raw footage frames (768 x 768 pixels sized) from 12-bit to 8-bit grey scale, then stacked over three channels to produce RGB images. An example of the resulting images is shown in Figure 4-2 and will be referred to as *rgb* type images. Another set of images for each experimental run were also produced in which three consecutive 8-bit grey scale frames were stacked over the three channels (i.e., Red = frame – 2, Green = frame – 1, Blue = frame). An example of which is shown in Figure 4-2 and will be referred to as *rgb_motion* type images. Though the pretrained CNN used has not been trained on images of this nature, it was trialed to see if the features extracted improved the trained model performance.

Figure 4-2*Image Types and Crop Positions*

Note. Three channel rgb image (left) and rgb_motion image (centre). Crop positions trialled for classification analysis (right).

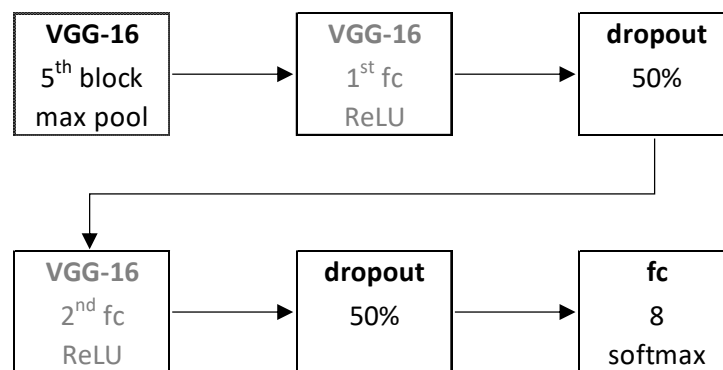
To construct the training dataset, all the experimental underflow footage was reviewed and only experimental runs with extended periods of relatively stable underflow profile were selected for model training. From these runs, 8000 consecutive images were used as the training dataset, and 2000 consecutive images from later in the footage were used as the validation dataset. Appendix E details the specific frames used in the various datasets. The images were cropped to 224 x 224 pixels, to match the expected input dimension of the pretrained CNN. The cropped images were centred horizontally based on the apex opening, with Position A and Position C trialled as two distinctly different crop positions (Figure 4-2). Concerns with droplets forming on the apex structure potentially resulting in a consistent presence in some footage led to the testing of the slightly lower Position B.

For training the classifier, DIGITS' (NVIDIA Corporation, 2014) Caffe (Jia et al., 2014) implementation of VGG-16 (Simonyan & Zisserman, 2014), pretrained on ILSVRC2012 dataset was used. As the cropped image sizes were already 224 x 224 pixels, only mean pixel subtraction was performed as preprocessing. The pretrained weights of the convolutional layers were used, and the layers frozen (lr_mult: 0). The first two fully connected layers were used with their pretrained weights but fine-tuned during training (lr_mult: 1.0). VGG-16's original final fully connected layer was

replaced with an eight node fully connected layer, given the eight class classification task, with the weights updated more aggressively (lr_mult: 10.0). The key trainable portion of the architecture described is shown in Figure 4-3. The optimiser used was stochastic gradient decent with a base learning rate of 1×10^{-4} , decaying by a factor of 0.1 every 10 epochs, with a training batch size of 32 used throughout, and trained for 30 epochs.

Figure 4-3

Hydrocyclone Underflow Footage Classification Network Architecture



Note. Shown from the VGG-16 (Simonyan & Zisserman, 2014) fifth block max pooling layer, and with layers using pretrained weights for fine-tuning shaded grey. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

4.1.2.2 Testing and Discussion

The trained model's performance on the different image types and crop positions for the training dataset, validation dataset, and test dataset (i.e., remaining frames from each relevant experimental run) are summarised in Table 4-1, Table 4-2 and Table 4-3 respectively. For both image types, Position A shows better training and validation dataset accuracies than Position C. Hydrocyclone design and operating conditions dictate the underflow profile, as indicated by spray angle investigations such as van Deventer et al. (2003) and Dubey et al. (2016). Also, the further from the

apex the more external forces (such as gravity) would influence the underflow profile, along with ultimately the breaking up of the continuous sheet-like flow (van Deventer et al., 2003). Therefore, it is reasonable that Position C with a more dispersed underflow appearance, less influenced by the hydrocyclone discharge conditions, and without the edges of the underflow visible would have less meaningful features present to correlate to the current operating condition than Position A.

Table 4-1

Training Dataset Footage Classification Model Accuracy

Image Type	Position	Overall Accuracy (%)	Class Accuracy (%)
rgb	A	99.97	99.89–100.00
	B	99.28	96.99–100.00
	C	94.56	89.88–99.06
rgb_motion	A	99.99	99.96–100.00
	B	99.40	98.03–100.00
	C	92.64	88.36–97.01

Note. Overall accuracy and range of per class accuracies for varying image types and crop positions trialed.

Table 4-2

Validation Dataset Footage Classification Model Accuracy

Image Type	Position	Overall Accuracy (%)	Class Accuracy (%)
rgb	A	93.81	59.75–99.80
	B	95.90	93.10–99.30
	C	75.82	56.40–92.85
rgb_motion	A	94.02	57.90–99.85
	B	96.52	91.75–99.65
	C	67.04	46.65–82.80

Note. Overall accuracy and range of per class accuracies for varying image types and crop positions trialed.

Table 4-3*Test Dataset Footage Classification Model Accuracy*

Image Type	Position	Overall Accuracy (%)	Class Accuracy (%)
rgb	A	85.49	34.67–99.55
	B	95.09	91.17–99.40
	C	-	-
rgb_motion	A	85.78	34.87–99.65
	B	95.49	90.41–99.43
	C	-	-

Note. Overall accuracy and range of per class accuracies for varying image types and crop positions trialled.

Comparing the results of Position A and Position B, shows both have a strong fit to the training dataset. The overall validation dataset accuracy of Position B is slightly higher and more importantly the per class accuracies are all >90%, while for crop Position A there are individual class accuracies <60%. When the remaining frames from each relevant experimental run are used as a test dataset, unseen during the training process, the greater performance of the model produced from Position B becomes more evident. The overall accuracy and per class accuracy range for Position B is similar to the validation dataset, while for Position A the performance is lower than the validation dataset thus showing a poorer performance on new data. Though there is only a small difference between the images of Position A and Position B, it appears that the shift to include more of the wider highly textured band lower in the underflow stream and less of the narrower underflow band near the apex was sufficient to improve class discernment.

Overall, the models based on rgb_motion images did not appear to give a significantly better accuracy than those based on rgb images. This is not to say that including temporal information into CNN based models will not yield improvements, but that this pretrained model constructed to classify colour images of objects was

not able to adequately use this information to provide any further footage classification benefit. Given that through fine-tuning VGG-16 using crop position B images features were able to be extracted and used to classify these experimental runs successfully (>90% per class accuracy), the next stage of the investigation explored if the extracted features can be used to estimate particle size information of the hydrocyclone underflow stream.

4.1.3 Particle Size Regression

4.1.3.1 Model Development

The images from rgb image type and crop Position B, as in Section 4.1.2.1, made up the training and validation dataset for the regression model, and were paired with the measured hydrocyclone underflow P80 for the corresponding experimental run. Given the three channels of the images are identical, and that they differ from the ImageNet (Deng et al., 2009) dataset used in pretraining model's available in Tensorflow's (Abadi et al., 2015) implementation of Keras (Chollet & others, 2015a), the only preprocessing performed was training dataset mean pixel intensity subtraction from all images.

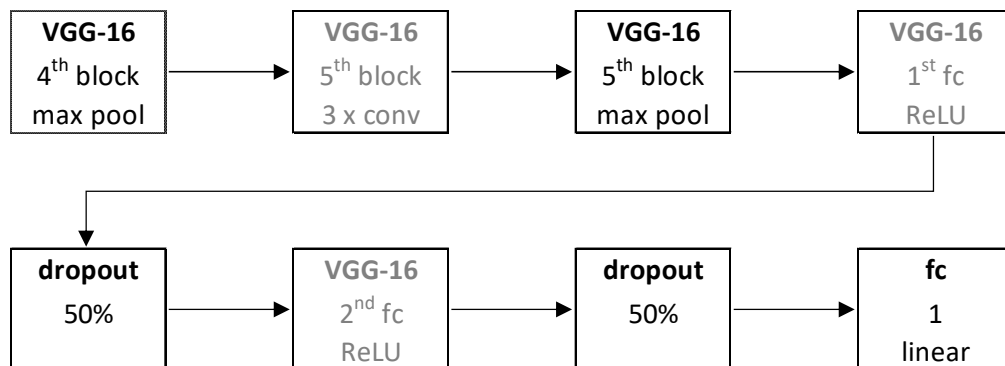
The mean centred image data was then passed through the Keras implementation of VGG-16 (Simonyan & Zisserman, 2014). Initially, on top of the output of the second fully connected layer of VGG-16 a 50% dropout layer (Srivastava et al., 2014) was added, followed by a single node fully connected layer to output the predicted hydrocyclone underflow P80. A training batch size of 32 was used throughout, and with mean absolute error (MAE) as the loss function given that it is less influenced by outliers than MSE (Draper & Smith, 1998). While periods of relatively stable underflow profile make up the training and validation dataset the intent is for any frames resulting from instability or containing atypical features, and thus not representative of the typical underflow profile, to have a lower influence on model development. With only the weights of the new output layer trainable, using Adam (Kingma & Ba, 2014) optimizer with a learning rate = 1×10^{-3} , the network was then trained for 3 epochs. The brief training of the new output layer was performed so that the weights are more relevant to the task prior to fine-tuning deeper in the

network, rather than attempting fine-tuning with an output layer of randomly initialised weights (Chollet, 2018).

Next the parameters of the VGG-16 network up to the max pooling layer following the fourth convolution block were frozen. The network's parameters from the fifth convolutional block through to the newly added output layer were kept trainable for fine-tuning. Dropout layers with 50% probability were also included between the first and second fully connected layers, along with between the second fully connected layer and output layer. The key trainable portion of the architecture described is shown in Figure 4-4. Adam optimizer was again used, though with a lower learning rate of 1×10^{-5} and decay rate of 1×10^{-6} . The network was then trained for 10 epochs, with a final validation dataset MAE of $6.2 \mu\text{m}$.

Figure 4-4

Hydrocyclone Underflow Particle Size Regression Network Architecture



Note. Shown from the VGG-16 (Simonyan & Zisserman, 2014) fourth block max pooling layer, with the number of convolutional (conv) layers in each block, and the layers using pretrained weights for fine-tuning shaded grey. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

4.1.3.2 Testing and Discussion

Table 4-4 includes simple statistics for the validation dataset's predicted underflow P80, with Figure 4-5 showing a plot of the predicted and measured underflow P80.

While the model provides an acceptable fit to the validation data this was expected considering the data is interrogated during the training process.

Table 4-4

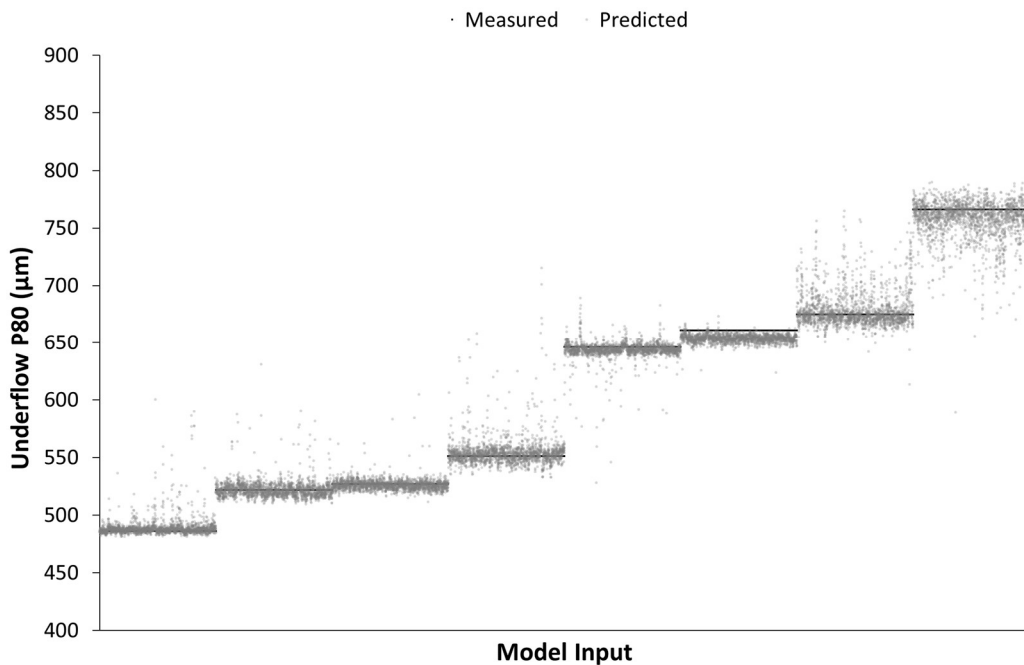
Laboratory Validation Dataset Hydrocyclone Underflow P80 Performance Statistics

Measured P80 (μm)		486	522	527	551	646	661	675	766
Predicted P80 (μm)	Mean	489	522	526	554	644	654	678	758
	SD	8	7	5	13	8	3	15	16

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-5

Laboratory Validation Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each footage period, with the periods then sorted by increasing measured underflow P80.

A more meaningful gauge of the model's performance would be through making predictions on footage of experimental runs not included in the training and validation datasets. Thus, a test dataset was constructed from the remaining experimental runs that were previously deemed to lack a sufficient period of relatively stable underflow profile to be included in the training dataset. From each of these experimental runs 2000 consecutive images of relatively stable underflow profile were selected. Appendix F details the specific frames used in the test dataset.

Table 4-5 includes simple statistics for the test dataset's predicted underflow P80, with Figure 4-6 showing a plot of the predicted and measured underflow P80. While the mean predicted underflow P80 from the coarser feed examples are not as accurate as those achieved from the finer feed examples, the predictions are reasonable in that the footage from the distinct fine and coarse feed runs are correctly allocated as such. In terms of precision, there are several cases in which the predicted underflow P80 standard deviation is greater than that observed in the validation dataset, though given these frames were from experimental runs not used in training and noted for having greater periods of instability this is to be expected.

Table 4-5

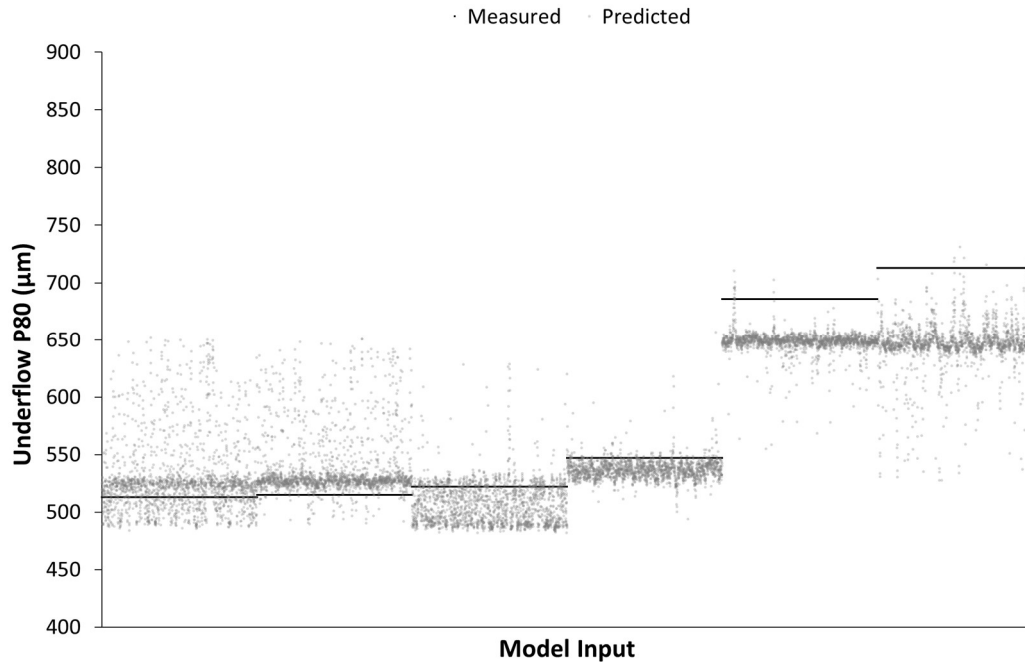
Laboratory Test Dataset Hydrocyclone Underflow P80 Performance Statistics

Measured P80 (μm)		513	515	522	547	686	713
Predicted P80 (μm)	Mean	528	537	507	537	649	647
	SD	31	27	18	9	9	20

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-6

Laboratory Test Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each footage period, with the periods then sorted by increasing measured underflow P80.

While an acceptable number of image frames were available for model training, the limited number of useable experimental runs and thus particle size examples also limit what can be interpreted from the model's performance. Comparing the validation and test dataset's performance it is not certain that the relationship between image features established is capable of being extended to unseen conditions. Though the results were sufficient to warrant further investigation under more industrially applicable conditions and using a more widely available lower framerate colour video camera. With a fixed camera setup, a similar filming and sampling process over a period of circuit operation would build a large dataset under varying operational and environmental conditions. A CNN regression modelling trial would probe if from images alone an acceptable underflow particle size prediction could be produced.

Considering the improvements yielded from incorporating unconventional features into empirical models of hydrocyclones (as discussed in Section 2.2) and Uahengo's (2014) recommendation to explore incorporating additional measurable parameters into image-based modelling (as discussed in Section 2.4.4), the combination of CNN extracted features with other available hydrocyclone circuit data into a regression model was of interest. Lin and Jørgensen (2011) is one such industrially relevant example, where features were extracted from footage of an operating cement kiln using multivariate image analysis and combined with measured PVs in a regression model to predict nitrogen oxides emissions. This was found to have improved performance compared to soft sensors developed on conventional process measurements alone.

4.2 Industrial Hydrocyclone


This investigation builds upon the proof-of-concept work undertaken in Section 4.1, extending it to data collection under industrial conditions and exploring the incorporation of additional sensor information into the underflow particle size prediction modelling. All videos, images, data, and information from IGO's Nova mine site referred to throughout Section 4.2 are captured by the dataset IGO Nova (2020) and used with permission.

4.2.1 Data Collection

This experiment was conducted on the primary hydrocyclone cluster of IGO's Nova nickel-copper sulphide processing plant. Ambient lighting behaviour throughout the day was a key consideration in selecting the location for filming. The underflow stream of the hydrocyclone selected for monitoring was illuminated by indirect sunlight for most of the day (along with the additional light source applied for the experiment). Only later in the day, and when not obscured by clouds, could direct sunlight result in overexposed regions in the frame as illustrated in Figure 4-7. In practice, enclosing the area to control for external light would be desirable, such that the applied light is the sole source of illumination for the underflow and thus minimising a source of visual variability.

Figure 4-7

Examples of Illumination Variation Throughout a Single Day



Note. Date stamp has been redacted. Adapted from *Nova Mine Site Data*, by IGO Nova, 2020. Copyright 2020 by IGO. Adapted with permission.

A Panasonic Lumix G DC-G9GN-K camera fitted with a Panasonic Lumix G X Vario 12–35 mm / F2.8 II H-HSA12035E lens was used to film the hydrocyclone underflow. A Promaster HGX Prime 58 mm UV Filter (code 6711) was also attached to the front of the lens for additional lens protection and ease of cleaning when required. A wired remote shutter release was attached to the camera, so that starting and stopping the camera filming did not require any of the buttons on the camera to be touched, thus minimising the introduction of any movement/vibration in this process.

A scaffold platform was erected on the underflow launder level grid mesh for camera mounting. The camera was attached to a tripod (with legs at minimum extension) which was fixed to the scaffold, and with a sandbag suspended from the centre pole. The intention when planning this setup was to hold the camera in a fixed

position for the duration of the experiment and attempt to minimise the effect of vibration (mainly that caused by wind and operating plant equipment).

A LED light source was directed at the underflow discharge stream and mounted in a position where it did not interfere with the camera's view of the underflow discharge stream. Figure 4-8 shows the resulting experimental setup, with the distance from the outer surface of the spigot housing to the LED light being 0.35 m and to the focal plane mark on the camera body being 1.37 m.

Figure 4-8

Experimental Setup at IGO Nova Primary Hydrocyclone Cluster



Note. Adapted from *Nova Mine Site Data*, by IGO Nova, 2020. Copyright 2020 by IGO. Adapted with permission.

The relevant camera settings for filming the underflow were as follows

- Recording Format = MP4
- Resolution = 4K
- Bite Rate = 150 Mbps
- Frame Rate = 60 FPS
- Luminance Level = 0–255

- Aperture = f/2.8
- Shutter Speed = 1/3200 s
- ISO Sensitivity = ISO3200
- White Balance = WB K Set 6500K (to match LED colour temperature)
- Photo Style = Natural

All camera and lens image stabiliser options were turned off. Manual focussing was performed, aided by the camera's MF Assist with focus peaking (HIGH Detect Level) to guide focusing on the underflow discharge stream. Once the focus was deemed satisfactory, the focus ring was then locked, and lens position resume activated to ensure a consistent focus throughout the experiment.

Of note, with a fast shutter speed used to minimise motion blur, and given the camera's electronic rolling shutter, the nature of the additional light source selected was an important consideration. Flickering results from variation in light source luminosity and while ideally an LED light source would be effectively flicker free, factors such as LED quality, type of DC adapter, and use of dimming can lead to flickering of varying degrees (Pueo, 2016). Figure 4-9 illustrates the banded appearance of footage that can occur with the camera settings stated previously and a flickering LED light source. This issue was ultimately avoided by using a 50W 5000 lumen rechargeable LED light (Detroit DET50WRELED) as the additional light source.

Figure 4-9

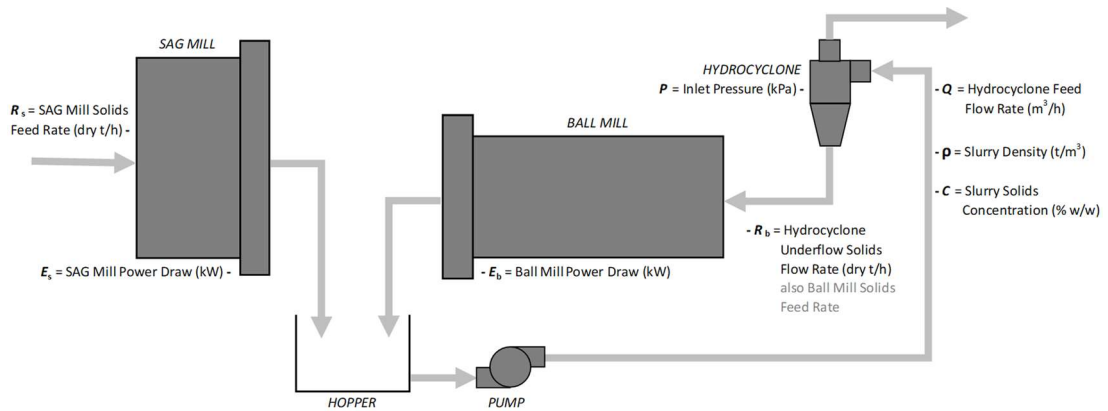
Banded Footage Example Frames



Note. Shown are three consecutive frames using the camera settings as previously detailed and a LED light source which results in frames with a banded appearance.

Prior to filming each data collection event the light was cleaned and camera lens protector removed. The remote shutter release was then used to start recording, with a minimum of 32 s worth of footage collected. A sampling cut of the entire underflow stream was then taken using a metallurgical sample cutter for representative sampling. The underflow samples were pressed, oven dried, and riffle split sub-sampled for particle size analysis via dry sieving, which was performed as soon as practical to minimize the risk of sample loss. No external assistance was sought, with the entire experimental procedure undertaken by one person to minimise sources of variability and error.

Available sensors readings and subsequent calculated variables from the comminution circuit were also provided at 1 s intervals for each data collection event. Figure 4-10 shows a simplified representation of IGO Nova's primary grinding circuit with the location of sensors readings and calculated variables used as model inputs in this chapter indicated. While an attempt was made to synchronize the camera's timestamp with the SCADA system, a confirmation check at the end of the experimental period showed the camera's timestamp to be approximately 3 s behind the time displayed on SCADA. This offset was assumed to be consistent throughout the experimental period when allocating sensor data to the time-stamped camera footage.

Figure 4-10*IGO Nova Simplified Primary Grinding Circuit Diagram*

Note. IGO Nova's primary grinding circuit with open circuit semi-autogenous grinding (SAG) mill and closed circuit ball mill with hydrocyclone cluster. Location of sensors readings and calculated variables indicated.

The data collection events that formed the training dataset (Train) were collected over four consecutive days, and those for the validation dataset (Val) were collected over the day following this period. These data collection events occurred approximately hourly from 7:00AM to 4:00PM each day. All data collection events were included in the training and validation datasets if the footage did not contain any bright over exposed areas within the hydrocyclone underflow discharge region. This resulted in the training dataset being made up of 35 data collection events, and the validation dataset made up of 9 data collection events.

The test datasets were produced from data collection events the day after the validation dataset was collected (Test NA) and 2 days prior to the first day of training dataset collection (Test NB). Of note these data collection events were not taken specifically at hourly intervals to give variation in the sampling times. The Test NB footage was also collected during a period in which adjustments were made to the camera's focus, so though the appearance was visually similar to subsequent footage, it was prior to the focus being set and locked from the training dataset collection onwards. Footage containing over exposed regions in the hydrocyclone

underflow discharge region were also included in the test datasets to gauge the trained model's performance under these conditions. The Test NA dataset was made up of seven data collection events, and the Test NB dataset made up of five data collection events.


The decision to construct the datasets using blocks of daily data was made to better align the training process to the nature of an applied sensor. In practice training would be performed on an existing collection of data and then the trained model would be expected to accurately predict future hydrocyclone underflow P80 values. Chollet (2018) highlights issues related to dataset construction in terms of leakage of information and of importance to this modelling task the consideration of leakage of temporal information. Issues related to the leakage of information between the training, validation, and testing dataset were minimised as each entire data collection event is only used in one specific dataset, instead of for example the random splitting of frames and relevant sensor information between the datasets. Issues related to the leakage of temporal information is also minimised by having the training and validation datasets split as sequential daily blocks. The test dataset is then constructed from daily blocks outside this period, rather than for example randomly allocating each data collection event into one of the datasets.

Figure 4-11 presents the measured sample hydrocyclone underflow P80 of all the data collection events used to form the relevant datasets, in the order that they were collected and showing the variability of underflow P80 over time. Figure 4-11 also indicates periods where notable operational changes have occurred. Typical operation of the primary hydrocyclone cluster during the experimental period had six hydrocyclones operating, except for the period indicated on day four in which five hydrocyclones were operating (marked as 5 CYC in Figure 4-11). Also, on days four and five the processing plant was being fed from the emergency run-of-mine (EROM) ore stockpile, which is known to be a typically coarser feed source. Table 4-6 shows the overall feed ore blend for each day during the experimental period. Figure 4-12 presents the measured sample hydrocyclone underflow P80 of all the data collection events plotted against the sampling time. This shows no clear relationship between these variables, thus the degree of illumination depending on the time of day or any

operational processes undertaken at a certain time of day are unlikely to have an influence on the subsequent model developed.

Figure 4-11

Industrial Dataset's Hydrocyclone Underflow P80 in Order Sampled

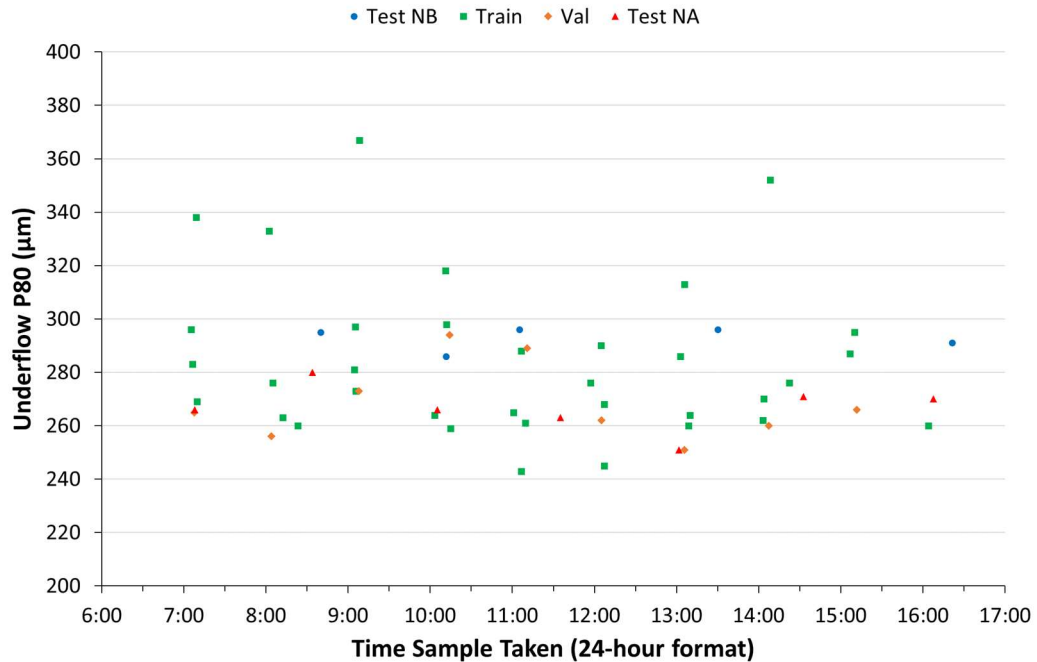


Note. Presents IGO Nova (2020) hydrocyclone underflow P80 for all datasets in the order samples were taken, with notable changes in operating conditions of emergency run-of-mine (EROM) feed ore and five hydrocyclone (5 CYC) highlighted.

Table 4-6*Daily Feed Ore Blend*

Dataset	Day	Feed Ore Blend (% w/w)					
		EROM	HG	MG	LG	VLG	PASTE
Test NB	1	0	32	21	34	8	5
Train	3	0	40	19	11	21	9
Train	4	100	0	0	0	0	0
Train	5	100	0	0	0	0	0
Train	6	0	35	27	20	9	10
Val	7	0	35	27	20	9	10
Test NA	8	0	35	27	20	9	10

Note. Presents daily feed ore blend from IGO Nova (2020) during the experimental period. Notation as per the Nova mine site designations; stockpile (EROM), high grade (HG), medium grade (MG), low grade (LG), very low grade (VLG), and paste backfill contaminated (PASTE).

Figure 4-12*Industrial Hydrocyclone Underflow P80 Against Sampling Time*

Note. Presents IGO Nova (2020) hydrocyclone underflow P80 for all datasets plotted against the time of sampling.

4.2.2 Image Only Modelling

4.2.2.1 Model Development


The 60 FPS camera frame rate setting was the highest available at 4K resolution (3840 x 2160 pixels), thus giving the most example images per second for training and testing. This setting yields a frame rate of 59.94 FPS in practice, resulting in certain 1 s periods in which there are only 59 frames with that timestamp rather than the expected 60 frames. For simplicity, in dataset construction from the first timestamp for the period of interest each 60 frame block was assumed to belong to 1 s worth of footage. This resulted in some frames being slightly out of sync when subsequently paired with sensor data, given the practical frame rate issue. Considering the brief time periods used (i.e., 32 s and thus 1920 frames for each data collection event in the training and validation datasets) and in the context of the

overall difficulty in synchronizing the timestamp of a battery powered camera in the field with the SCADA system, this was deemed acceptable. For relevant data in subsequent investigations in which image frames are paired with other sensor data 60 images were then matched with 1 s worth of sensor data.

The OpenCV (Bradski, 2000) Python library (opencv-python package; Heinisuo et al., 2019) was used to extract the frames from the video footage. The Pillow (Clark & contributors, 2010) Python library was then used to extract an 896 x 896 pixels crop region from a fixed location in each frame. This crop position is below the shadow cast by the spigot housing onto the underflow stream and roughly centres the stream, as illustrated in Figure 4-13. This being a visually similar position to crop position B shown in Figure 4-2 and used in Section 4.1.

Figure 4-13

Example Video Frame Showing the Crop Location



Note. Date stamp has been redacted. Adapted from *Nova Mine Site Data*, by IGO Nova, 2020. Copyright 2020 by IGO. Adapted with permission.

The cropped frames were then resized to 224 x 224 pixels and the required input preprocessing performed as expected by the pretrained CNN. In this case, the bicubic interpolation method was used rather than the Keras (Chollet & others, 2015a) default nearest interpolation method, as a trade-off between retaining image quality and speed of processing images (Clark & contributors, n.d.).

The model was then trained using the same method applied for the laboratory hydrocyclone images in Section 4.1.3. The preprocessed images were passed through Keras' (Chollet & others, 2015a) pretrained implementation of VGG-16 (Simonyan & Zisserman, 2014). Initially, on top of the output of the second fully connected layer of VGG-16 a 50% dropout layer (Srivastava et al., 2014) was added, followed by a single node fully connected layer to output the predicted hydrocyclone underflow P80. MAE was again used as the loss function, and a training batch size of 32 used throughout. With only the weights of the new output layer trainable, using Adam (Kingma & Ba, 2014) optimizer with a learning rate of 1×10^{-2} , the network was then trained for 3 epochs.

Next the parameters of the VGG-16 network up to the max pooling layer following the fourth convolution block were frozen. The network's parameters from the fifth convolutional block through to the newly added output layer were kept trainable for fine-tuning. Dropout layers with 50% probability were also present between the first and second fully connected layers, along with between the 2nd fully connected layer and output layer. Adam optimizer was again used, though with a lower learning rate of 1×10^{-4} and decay rate of 1×10^{-2} , and the model trained for 5 epochs. This network training method will be referred to as convtune.

4.2.2.2 Testing

The final validation set's MAE was 13.1 μm , with Table 4-7 showing simple statistics for the validation dataset's predicted underflow P80, and Figure 4-14 showing a plot of the predicted and measured underflow P80. As the validation dataset does not include examples from when the EROM stockpile was used to feed the processing plant, Figure 4-15 showing a plot of the predicted and measured underflow P80 for the training dataset has also been included, with a MAE of 7.6 μm for the training set noted.

Table 4-7


Conv tune Model Validation Dataset Hydrocyclone Underflow P80 statistics

Measured P80 (μm)		265	256	273	294	289	262	251	260	266
Predicted	Mean	275	265	270	265	272	270	263	266	284
P80 (μm)	SD	10	8	7	6	7	5	4	6	7

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-14


Conv tune Model Validation Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-15

ConvTune Model Training Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each sampling period, with the periods then presented in the order samples were taken.

The test datasets were then passed through the trained convtune model. Test NA resulted in a MAE of 8.4 µm, with Table 4-8 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-16 showing a plot of the predicted and measured underflow P80. Test NB resulted in a MAE of 14.6 µm, with Table 4-9 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-17 showing a plot of the predicted and measured underflow P80.

Table 4-8


Conv tune Model Test NA Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		266	280	266	263	251	271	270
Predicted P80 (μm)	Mean	269	281	273	270	265	275	277
	SD	10	9	7	8	4	5	5

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-16

Conv tune Model Test NA Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Table 4-9


Conv tune Model Test NB Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		295	286	296	296	291
Predicted	Mean	274	273	272	286	289
P80 (μm)	SD	5	6	7	6	7

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-17

Conv tune Model Test NB Dataset Measured and Predicted Underflow P80



Note. Model input images ordered by time for each sampling period, with the periods then presented in the order samples were taken.

4.2.3 Sensor Only Modelling: Hydrocyclone


4.2.3.1 Model Development

The first model incorporating sensor data only included available variables directly related to the hydrocyclone, thus trialling some of the variables used in the cut size ANN modelling undertaken by van Loggenberg et al. (2016). The variables used in the first model were

- P = Hydrocyclone inlet pressure (kPa)
- Q = Hydrocyclone feed flow rate (m^3/h)
- ρ = Hydrocyclone feed slurry density (t/m^3)
- C = Hydrocyclone feed solids concentration (% w/w)

The number of operating hydrocyclones was not explicitly included as inputs in the model, but the relationship between pressure and feed flow rate could provide an implicit identification of these distinct operating conditions encountered in the training dataset if found to be meaningful in the training process.

There were 32 s worth of sensor readings used for each data collection event in the training and validation datasets. Each input variable for the model was first standardised by subtracting the variable's training dataset mean and then dividing by the variable's training dataset standard deviation. A simple ANN was then constructed with Keras (Chollet & others, 2015a) by adding a three node fully connected hidden layer with ReLU activation, followed by a single node fully connected layer to output the predicted hydrocyclone underflow P80, as illustrated in Figure 4-18. MAE was again used as the loss function, and with a training batch size of 32 used. Adam (Kingma & Ba, 2014) optimizer with a learning rate of 1×10^{-1} and decay rate of 1×10^{-4} was used to train the network for 14 epochs. This network training method will be referred to as *cycsensor*.

Figure 4-18*Neural Network Architecture for Cycsensor Model*

Note. Includes the number of nodes and activation for fully connected (fc) layers.

4.2.3.2 Testing

The final validation set's MAE was 10.2 μm , with Table 4-10 showing simple statistics for the validation dataset's predicted underflow P80, and Figure 4-19 showing a plot of the predicted and measured underflow P80. The training set's MAE was 12.2 μm , with Figure 4-20 showing a plot of the predicted and measured underflow P80. The test datasets were then passed through the trained cycsensor model. Test NA resulted in a MAE of 37.8 μm , with Table 4-11 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-21 showing a plot of the predicted and measured underflow P80. Test NB resulted in a MAE of 16.5 μm , with Table 4-12 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-22 showing a plot of the predicted and measured underflow P80.

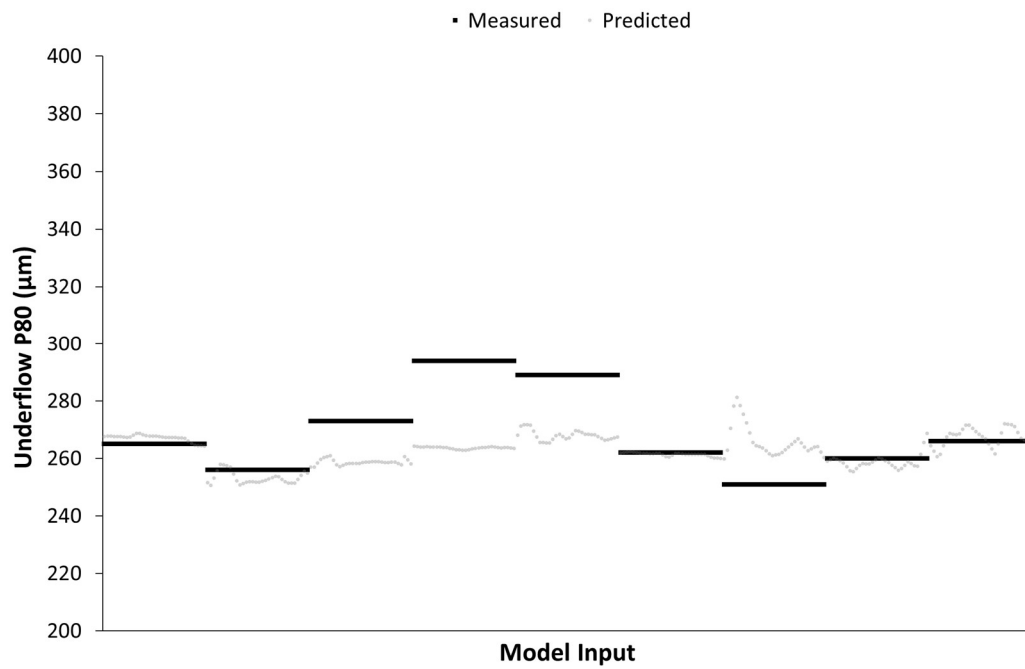
Table 4-10*Cycsensor Model Validation Dataset Hydrocyclone Underflow P80 Statistics*

Measured P80 (μm)		265	256	273	294	289	262	251	260	266
Predicted	Mean	267	253	259	264	268	261	266	259	267
P80 (μm)	SD	1	2	1	0	2	1	5	3	3

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-19


Cycsensor Model Validation Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-20

Cycsensor Model Training Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Table 4-11


Cycsensor Model Test NA Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (µm)		266	280	266	263	251	271	270
Predicted P80 (µm)	Mean	272	271	273	259	334	351	347
	SD	1	1	1	0	3	2	3

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-21

Cycsensor Model Test NA Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Table 4-12


Cycsensor Model Test NB Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (µm)		295	286	296	296	291
Predicted	Mean	283	288	272	314	317
P80 (µm)	SD	3	3	2	4	3

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-22

Cycsensor Model Test NB Dataset Measured and Predicted Underflow P80




Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.


Of note, there are distinctly higher underflow P80 predicted values for the final three Test NA data collection events (see Figure 4-21 and Table 4-11). Figure 4-23 shows the contribution of each component to the model's output underflow P80 value for each Test NA data collection events. Compared to the earlier events, during the final three events the third hidden node becomes active, as indicated by the increase in the bias contribution. The hydrocyclone feed slurry density also becomes a major positive contributor to the output value. Figure 4-24 shows the hydrocyclone feed slurry density readings for all the datasets. The final three Test NA data collection events' readings are low and outside that found in the training dataset. The current model is not able to successfully handle hydrocyclone feed slurry density readings outside the range used in training.

Figure 4-23

Cycsensor Model Components Contribution to Output Values for Test NA Dataset



Note. Shows model inputs variables hydrocyclone inlet pressure (P), feed flow rate (Q), feed slurry density (ρ) and feed solids concentration (C), along with Bias (as the sum of all layer biases), contribution to the predicted P80 output. Input data is ordered by time, with the corresponding measured P80 for the sampling period indicated.

Figure 4-24*Hydrocyclone Feed Slurry Density Readings for All Datasets*

Note. Presents IGO Nova (2020) hydrocyclone feed slurry density (ρ) readings for all datasets in the order samples were taken, with the final three Test NA samples' readings highlighted.

4.2.4 Sensor Only Modelling: Extended

4.2.4.1 Model Development

The second model incorporating sensor data only then included additional calculated and measured variables from the wider grinding circuit, taking inspiration from some of the variables also included in Zhang and Liang's (2016) closed circuit primary ball mill particle size soft sensor investigation. Considering the influence of hydrocyclone feed slurry density on the cycsensor model predictions when the sensor readings were outside that found in the training dataset, this variable was not included in the second sensor data only model. The variables used in the model were


- R_s = SAG solids feed rate (dry t/h)
- E_s = SAG mill power draw (kW)

- P = Hydrocyclone inlet pressure (kPa)
- Q = Hydrocyclone feed flow rate (m^3/h)
- C = Hydrocyclone feed solids concentration (% w/w)
- R_b = Hydrocyclone underflow solids flow rate (dry t/h),
also ball mill solids feed rate
- E_b = Ball mill power draw (kW)

As with the cysensor model, 32 s worth of sensor readings were used for each data collection event in the training and validation datasets. Each input variable for the model was first standardised by subtracting the variable's training dataset mean and then dividing by the variable's training dataset standard deviation. A simple ANN, with a larger capacity than that used in cysensor model, was then constructed with Keras (Chollet & others, 2015a). A 10 node fully connected hidden layer with ReLU activation was used, followed by a single node fully connected layer to output the predicted hydrocyclone underflow P80, as illustrated in Figure 4-25. MAE was again used as the loss function, and a with training batch size of 32 used. Adam (Kingma & Ba, 2014) optimizer with a learning rate of 1.0 and decay rate of 1×10^{-1} was used to train the network for 9 epochs. This network training method will be referred to as gcsensor.

Figure 4-25

Neural Network Architecture for Gcsensor Model



Note. Includes the number of nodes and activation for fully connected (fc) layers.

4.2.4.2 Testing

The final validation set's MAE was 12.2 μm , with Table 4-13 showing simple statistics for the validation dataset's predicted underflow P80, and Figure 4-26

showing a plot of the predicted and measured underflow P80. Of note, for the seventh validation dataset sampling event (measured underflow P80 of 251 μm) there is a distinct spike in the early predictions. This resulted from an increase in the hydrocyclone feed flow rate and was also present in the cycsensor prediction for the same sample (see Figure 4-19). In the gcsensor case the related hydrocyclone underflow solids flow rate also contributes along with the hydrocyclone feed flow rate, resulting in a more pronounced spike during this less stable period of operation. The training set's MAE was 11.1 μm , with Figure 4-27 showing a plot of the predicted and measured underflow P80.

Table 4-13


Gcsensor Model Validation Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		265	256	273	294	289	262	251	260	266
Predicted	Mean	280	265	280	284	276	265	292	266	264
P80 (μm)	SD	4	8	2	5	3	3	11	3	5

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-26


Gcsensor Model Validation Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-27

Gcsensor Model Training Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

The test datasets were then passed through the trained gcsensor model. Test NA resulted in a MAE of 37.8 µm, with Table 4-14 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-28 showing a plot of the predicted and measured underflow P80. For the Test NA dataset there are markedly higher underflow P80 predicted values for the fifth data collection event (where measured underflow P80 is 251 µm). Figure 4-29 shows the contribution of each component to the model's output underflow P80 value for each Test NA data collection events. For the fifth data collection event SAG mill solids feed rate is a major positive contributor to the output value, distinct from the minor contribution in the rest of the Test NA dataset. Figure 4-30 shows the SAG mill solids feed rate readings for all the datasets. The fifth Test NA data collection event's readings are high and outside that found in the training dataset.

Table 4-14


Gcsensor Model Test NA Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		266	280	266	263	251	271	270
Predicted	Mean	305	289	281	308	347	295	307
P80 (μm)	SD	7	7	8	12	6	6	3

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-28


Gcsensor Model Test NA Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-29


Gcsensor Model Components Contribution to Output Values for Test NA Dataset



Note. Shows model input variables SAG mill solids feed rate (R_s) and power draw (E_s), ball mill solids feed rate (R_b) and power draw (E_b), hydrocyclone inlet pressure (P), feed flow rate (Q) and feed solids concentration (C), along with Bias (as the sum of all layer biases), contribution to the predicted P80 output. Input data is ordered by time, with the corresponding measured P80 for the sampling period indicated.

Figure 4-30

SAG Mill Solids Feed Rate Readings for All Datasets



Note. Presents IGO Nova (2020) SAG mill solids feed rate (R_s) readings for all datasets in the order samples were taken, with the fifth Test NA sample's readings highlighted.

Test NB resulted in a MAE of 26.2 μm , with Table 4-15 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-31 showing a plot of the predicted and measured underflow P80. For the Test NB dataset there are markedly lower underflow P80 predicted values for the third data collection event (the first event where measured underflow P80 is 296 μm). Figure 4-32 shows the contribution of each component to the model's output underflow P80 value for each Test NB data collection events. Only five to seven hidden nodes were active for the third event, compared to eight to ten active hidden nodes for the other events in the Test NB dataset. The contributions of the solids feed rate for the SAG mill and ball mill along with the hydrocyclone feed flow rate were negligible or slightly negative, rather than the positive contributions shown in the other Test NB examples. While there was a larger positive contribution from the hydrocyclone inlet pressure there

was also a larger negative contribution from the SAG mill power draw (which was also mainly outside those found in the training dataset as shown in Figure 4-33).

Table 4-15


Gcsensor Model Test NB Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		295	286	296	296	291
Predicted P80 (μm)	Mean	304	299	231	318	312
	SD	8	7	9	9	11

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-31

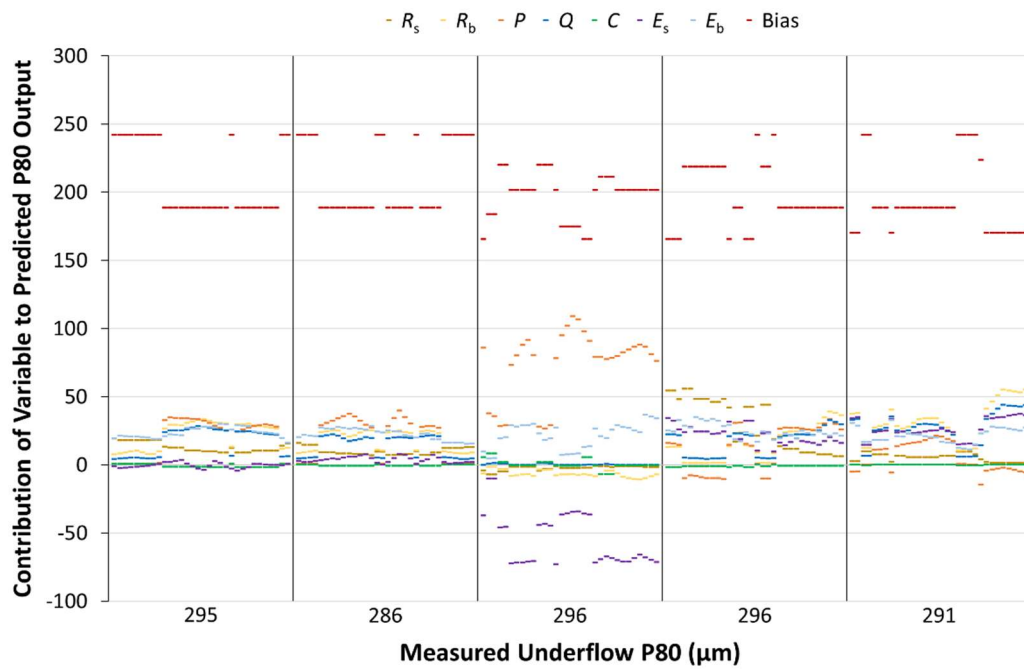
Gcsensor Model Test NB Dataset Measured and Predicted Underflow P80



Note. Model input data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-32


Gcsensor Model Components Contribution to Output Values for Test NB Dataset



Note. Shows model input variables SAG mill solids feed rate (R_s) and power draw (E_s), ball mill solids feed rate (R_b) and power draw (E_b), hydrocyclone inlet pressure (P), feed flow rate (Q) and feed solids concentration (C), along with Bias (as the sum of all layer biases), contribution to the predicted P80 output. Input data is ordered by time, with the corresponding measured P80 for the sampling period indicated.

Figure 4-33

SAG Mill Power Draw Readings for All Datasets



Note. Presents IGO Nova (2020) SAG mill power draw (E_s) readings for all datasets in the order samples were taken, with the third Test NB sample's readings highlighted.

4.2.5 Image and Sensor Modelling

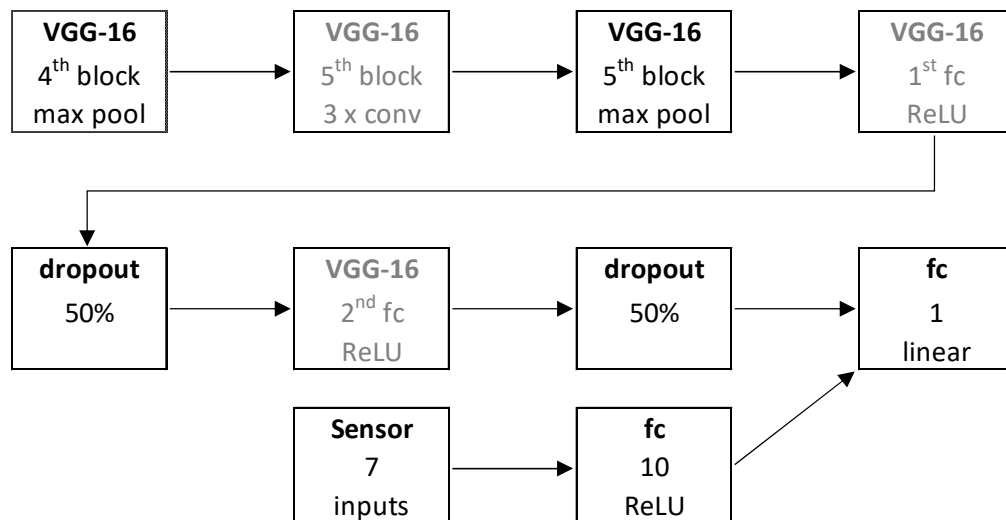
4.2.5.1 Model Development

To explore incorporating hydrocyclone underflow image data and comminution circuit sensor information into the regression modelling the image-only convtune model was combined with the larger sensor-only model gcsensor. The model architecture was effectively a concatenation of the convtune and gcsensor model outputs prior to the final layer. This was then connected to a single node fully connected layer to output the predicted hydrocyclone underflow P80, as illustrated in Figure 4-34. The input image and sensor information preprocessing were also the same as that described in the convtune and gcsensor methods, respectively. As mentioned in Section 4.2.2, given that there are 60 image frames per second and the sensor information was available at 1 s intervals the sensor readings for a given

timestamp are paired with each of the related 60 images for that 1 s period to form the input datasets.

Figure 4-34

Network Architecture for Conv-gc Model



Note. The image branch is shown from the VGG-16 (Simonyan & Zisserman, 2014) fourth block max pooling layer, with the number of convolutional (conv) layers in each block, and the layers using pretrained weights for fine-tuning shaded grey. Includes the number of nodes and activation for new fully connected (fc) layers and dropout layer rate.

For model training, MAE was again used as the loss function and a training batch size of 32 used throughout. With only the weights of the new 10 node fully connected hidden layer and new output layer trainable, and using Adam (Kingma & Ba, 2014) optimizer with a learning rate of 1×10^{-4} , the network was then initially trained for 2 epochs. Next the parameters of the VGG-16 network up to the max pooling layer following the fourth convolution block were frozen. The network's parameters from the fifth convolutional block through to the newly added output layer, along with the 10 node fully connected hidden layer, were kept trainable for fine-tuning. Adam optimizer was again used, though with a lower learning rate of

1×10^{-4} and decay rate of 1×10^{-6} , and the model trained for 8 epochs. This network training method will be referred to as conv-gc.

4.2.5.2 Testing

The final validation set's MAE was 11.0 μm , with Table 4-16 showing simple statistics for the validation dataset's predicted underflow P80, and Figure 4-35 showing a plot of the predicted and measured underflow P80. The training set's MAE was 1.2 μm , with Figure 4-36 showing a plot of the predicted and measured underflow P80. The test datasets were then passed through the trained conv-gc model. Test NA resulted in a MAE of 10.1 μm , with Table 4-17 showing simple statistics for the dataset's predicted underflow P80, and Figure 4-37 showing a plot of the predicted and measured underflow P80. Test NB resulted in a MAE of 18.6 μm , with Table 4-18 showing simple statistics for the dataset's predicted underflow P80 and Figure 4-38 showing a plot of the predicted and measured underflow P80.

Table 4-16

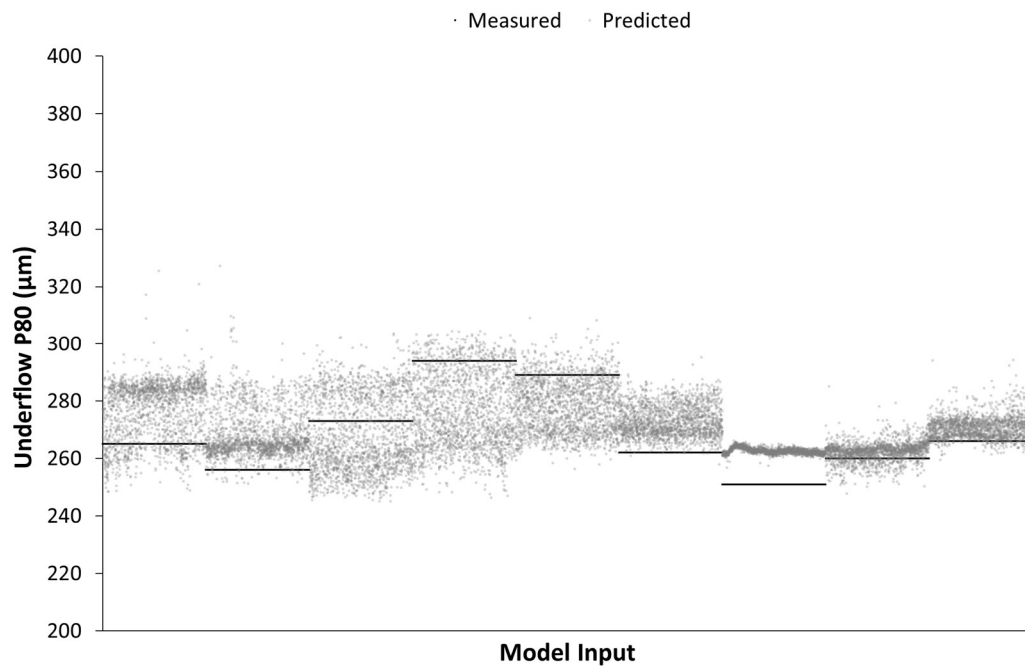
Conv-gc Model Validation Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (μm)		265	256	273	294	289	262	251	260	266
Predicted P80 (μm)	Mean	277	268	268	276	278	273	263	262	271
	SD	10	8	13	13	9	5	1	3	4

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-35


Conv-gc Model Validation Dataset Measured and Predicted Underflow P80



Note. Model input images/data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Figure 4-36

Conv-gc Model Training Dataset Measured and Predicted Underflow P80



Note. Model input images/data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Table 4-17


Conv-gc Model Test NA Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (µm)		266	280	266	263	251	271	270
Predicted P80 (µm)	Mean	265	279	265	276	267	271	277
	SD	9	15	9	11	2	5	6

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-37

Conv-gc Model Test NA Dataset Measured and Predicted Underflow P80



Note. Model input images/data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

Table 4-18


Conv-gc Model Test NB Dataset Hydrocyclone Underflow P80 Statistics

Measured P80 (µm)		295	286	296	296	291
Predicted P80 (µm)	Mean	286	272	268	278	268
	SD	6	5	8	7	2

Note. Predicted underflow P80 mean and population standard deviation (SD) shown against the measured underflow P80 for each data collection event.

Figure 4-38

Conv-gc Model Test NB Dataset Measured and Predicted Underflow P80



Note. Model input images/data ordered by time for each sampling period, with the periods then presented in the order samples were taken.

4.2.6 Discussion

During model training, hyperparameters (e.g., learning rate, decay rate, and number of epochs) were tuned with a focus on minimising the validation dataset MAE whilst also ensuring the training dataset MAE was reasonable. Table 4-19 summarises the trained model's MAE for the training, validation and test datasets for all models covered in Section 4.2. Each model's test dataset results are the best available indicator of their performance as they consist of events that are outside those used in the training process.

As noted in Table 4-6, the ore blend is the same during the final day of the training dataset and the validation dataset, which may lead to bias in the trained model towards this condition. This ore blend is also the same for the Test NA dataset, so it adds importance to the performance on the Test NB dataset which is from a

prior day and with a different ore blend. Again, it needs to be acknowledged that the potential different focus position during the filming of the Test NB dataset introduces another point of difference which may also affect an image-based model's performance.

Table 4-19

Trained Models' Mean Absolute Error for All Datasets

Inputs	Dataset MAE (μm)				
	Model	Training	Validation	Test NA	Test NB
Image Only	convtune	7.6	13.1	8.4	14.6
Sensor Only	cycsensor	12.2	10.2	37.8	16.5
Sensor Only	gcsensor	11.1	12.2	37.8	26.2
Image & Sensor	conv-gc	1.2	11.0	10.1	18.6

A more fine-grained comparison of the models' performance on the testing datasets can be found in Table 4-20 for Test NA and Table 4-21 for Test NB. These show the difference between the measured and average predicted hydrocyclone underflow P80 per data collection event for each model. The models developed solely on sensor data do not extend well to operating conditions outside that experience in training, with the MAE for the test datasets greater than those from the training and validation datasets. The most significant deviations from the measured values are when sensor readings are outside those found in the training range, as noted in Section 4.2.3.2 and Section 4.2.4.2. This highlights the requirement for the training and validation dataset to cover the full range of conditions to be encountered, and demonstrates the effect a single variable can have on the model's output. It would therefore be beneficial to only include key variables required to produce a meaningful predictive model. The risk of having excessive minor variables included is that if the underlying sensor or calculation were to fault or produce an atypical output this could then lead to an erroneous model output as well.

Table 4-20*Test NA Average Hydrocyclone Underflow P80 Difference for Data Collection Event*

Measured P80 (μm)		266	280	266	263	251	271	270
P80 Difference (μm)	convtune	3	1	7	7	14	4	7
	cycsensor	6	-9	7	-4	83	80	77
	gcsensor	39	9	15	45	96	24	37
	conv-gc	-1	-1	-1	13	16	0	7

Note. Hydrocyclone underflow P80 difference calculated as predicted – measured.

Table 4-21*Test NB Average Hydrocyclone Underflow P80 Difference for Data Collection Event*

Measured P80 (μm)		295	286	296	296	291
P80 Difference (μm)	convtune	-21	-13	-24	-10	-1
	cycsensor	-12	2	-24	18	26
	gcsensor	9	13	-65	22	21
	conv-gc	-9	-14	-28	-18	-23

Note. Hydrocyclone underflow P80 difference calculated as predicted – measured.

The models incorporating image data have resulted in testing dataset output MAEs more closely aligned with their relevant validation dataset MAE. These models also have a significantly higher capacity than the simpler sensor only models, which would be a contributing factor to their improved performance. The image and sensor combined conv-gc model resulted in a lower validation dataset MAE than the image only convtune model, but also showed a significantly lower training dataset MAE which may indicate overfitting to the training dataset. The conv-gc model also shows a Test NA dataset MAE similar to the validation dataset, though the Test NB dataset MAE is higher and a potential further indication of overfitting considering the ore feed similarities previously discussed. Ultimately, given that the conv-gc test dataset's MAE are also higher than that produced by the convtune model, the


inclusion of sensor data has not resulted in a pronounced improvement in model performance using the methods trialled.

To put the performance data for the models in context Figure 4-39 shows a box plot of the training dataset's measured hydrocyclone underflow P80 values (note the training dataset range of 124 μm). The image only convtune model presents the best testing performance in terms of MAE compared to the other models produced. Its highest data collection event standard deviation was 10 μm , occurring in both the validation dataset (see Table 4-7) and the Test NA dataset (see Table 4-8).

The convtune model's greatest absolute difference between predicted and measured underflow P80 was 29 μm where the validation measured P80 is 294 μm (see Table 4-7). For each test dataset the greatest absolute differences was 14 μm where the Test NA measured P80 is 251 μm (see Table 4-20) and 24 μm where the Test NB measured P80 is 296 μm (see Table 4-21). From the discussion of sensor data in Section 4.2.4.2 both the test data collection events noted are from periods of operations outside those experience in training, in these cases specifically the SAG mill operation. This may mean that operational conditions not experienced in training were not able to be interpreted by the image only model. Another issue may be that the model cannot assess more atypical or extreme underflow particle sizes, given the limited validation data range used to guide hyperparameter tuning during training. These more atypical conditions, such as occurs when ore is being fed from the EROM, were not able to be included in the validation and testing process given the limited time on site to collect data.

Figure 4-39

Box Plot of the Training Dataset Measured Hydrocyclone Underflow P80



Note. Constructed using the median exclusive method and with \times indicating mean.

One final point of interest is that the last data collection event for both test datasets had overexposed regions on the hydrocyclone underflow stream (Figure 4-40). Data collection events presenting over exposed regions were excluded from training and validation dataset construction, thus not encountered during model training. Though both models which use images were able to produce reasonable outputs from this footage. The mean P80 prediction difference from measured and prediction standard deviations for these data collection events were not outside those found in other examples from the test datasets. While ideally visual variability should be kept to a minimum and controlled where possible, further investigation into how robust the image-based models are to visual changes in conditions would be worthwhile. This would allow both assessment of the model's long-term stability and the requirement for external controls for lighting. If minimal additional lighting and structural modifications are required, then the more appealing and lower cost an image-based sensor would be.

Figure 4-40

Test Image Examples Exhibiting Over Exposed Regions



Test NA: Measured P80 = 270 μm Test NB: Measured P80 = 291 μm

Note. First frames of the test data collection events which exhibit overexposed regions (cropped as used for model input). Adapted from *Nova Mine Site Data*, by IGO Nova, 2020. Copyright 2020 by IGO. Adapted with permission.

To produce particle size inference models with neural networks (whether through available sensor information, image analysis or a combination of both) it is apparent that a sufficient quantity and quality of input data is required. Napier and Aldrich's (2017) conclusions from an investigation to develop an IsaMill regrind circuit product particle size soft sensor holds true not just to the sensor data case they explored but predictive models for mineral processing in general. They note that soft sensors would ideally require a check that the current conditions match those found in training. They also suggest that considering the variability encountered in mineral processing circuits (particularly in terms of feed) frequent recalibration when conditions change or a sufficiently long data collection period for training to try to capture enough data to compensate for process drift would be required. Further investigation into the potential for CNN based hydrocyclone underflow particle size estimation would benefit from a longer data collection period (for sufficient training

and testing data collection) along with further trials of model structures to gauge if a long term stable and usable sensor can be produced.

4.3 Summary

CNNs were investigated for their potential to extract meaningful features from hydrocyclone underflow video footage for particle size estimation. Pretrained VGG-16 (Simonyan & Zisserman, 2014) was used as a starting point in the model development, attempting to leverage the benefits of transfer learning.

Using high-speed laboratory hydrocyclone footage, initially different crop positions within the underflow stream were trialled as a classification task to determine where the most meaningful features for grouping footage could be found. A crop position near the hydrocyclone apex was found to perform best, with a >90% per class accuracy. This is consistent with van Deventer et al.'s (2003) videographic analysis, given that at this position the underflow profile is still highly influenced by the hydrocyclones operating conditions and the stream's edges are still present and defined.

Given the high-speed camera footage consisted of single channel grey scale images, while the pretrained CNN expects a three channel RGB image, the stacking of three consecutive frames was trialled. This was undertaken to explore if including temporal information in this way could be exploited to improve the classification performance. Under the conditions tested this did not result in a significant improvement in classification accuracy.

A regression model was also tested on the high-speed laboratory hydrocyclone footage to produce an underflow P80 particle size prediction. The best performing crop position from the classification task was used, and fine-tuning VGG-16 (Simonyan & Zisserman, 2014) from the fifth convolution block through to a new task specific output layer was performed. An acceptable model was able to be produced, as demonstrated by the validation dataset accuracy. When tested on footage from the remaining experimental runs not used in training, though the accuracy was lower discernment between the coarser and finer particle sizes was achieved, and thus warranted further investigation.

An investigation using a standard frame rate colour camera was then undertaken under industrially relevant conditions, with data collected from IGO Nova's primary hydrocyclone cluster. A similar experimental method as in the laboratory hydrocyclone investigation was performed, along with the collection of available online sensor readings for the grinding circuit. Simple ANN models were trialled using sensor data directly related to the hydrocyclone cluster as well as data from the wider grinding circuit. The subsequently trained models' performance suffered when tested on sensor readings outside those experience in training, highlighting the importance of having a representative training dataset.

Image-based models were also investigated, with a similar structure as that used in the laboratory hydrocyclone investigation along with trialling incorporation of grinding circuit sensor information. Both models resulted in improved testing dataset performance compared with the sensor-only models. The inclusion of sensor information into the CNN image model in this case did not result in improved performance compared to the image-based model alone. While the image-based model showed the best testing dataset performance, the long-term applicability and stability of such models would require more extensive investigation.

5 Reinforcement Learning Hydrocyclone Control

On reflecting on their investigation into milling circuit control incorporating RL Hallén et al. (2019) suggested that RL may find further success and industrial acceptance through control strategies combining it with PID controllers to improve robustness and performance guarantees, and with SP manipulation also the focus of Guo et al. (2019). This provides a good lead-in to the design and practical considerations envisaged when developing the investigation undertaken for this chapter. Ideally the RL controller's actions would adjust PID SPs to bring the system to meet targets without entering undesirable regions.

Other practical considerations are that in an industrial setting, actions taken on a live processing plant should be limited within a constrained region at each time step, aiding the system to change in a controlled manner. For example, with slurry being pump from a tank, the changes in density should be gradual. A sudden

aggressive increase in feed solids concentration can result in bogging of the pump, particularly in unmixed transfer hoppers. Another practical consideration is that sufficient sensor data is unlikely to be available to fully characterise the system (with sensor data itself also noisy). The sensor data availability and quality contribute to making live processing plant systems partially observable, which should be incorporated and managed in control design.

5.1 Experimental Design

5.1.1 Hydrocyclone Simulation

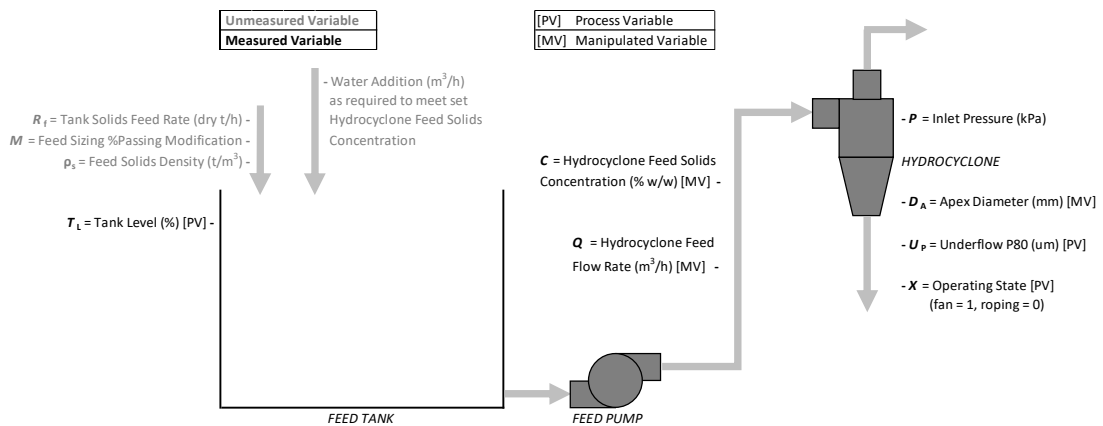
Moly-Cop Tools CycloSim (CycloSim_Single.xlsx) and the supplied base case example data (Sepúlveda, 2012) was used as the basis for the simulated hydrocyclone circuit. CycloSim is an Excel spreadsheet open circuit hydrocyclone simulator based on Gutiérrez and Sepúlveda's (1986; as cited in Sepúlveda, 2012) CIMM model. Moly-Cop Tools CycloSim was selected for the circuit simulation as it is supplied as a calibrated model, with the includes base case example data allowing it to be considered a physically plausible system. The other factor in simulation selection was its ability to be integrated with the RL controller. This was satisfied as permission was granted by Moly-Cop to implement the required equations and data in Python using the NumPy (Oliphant, 2006) and Math libraries. All CycloSim base case example hydrocyclone geometries, classifier constants, feed size distribution, ore density, and the number of operating cyclones as 10, were used as supplied unless otherwise stated. See Appendix G for the CycloSim Data_File spreadsheet containing the base case example data as supplied.

The simulated environment was designed to have the open circuit hydrocyclone cluster fed by a 500 m³ feed tank. A time step of 2 min was selected, thus it is assumed the input conditions are maintained for this period to allow the change in tank volume to be determined. The simulator also indicates the hydrocyclone underflow operating state of fan = 1 or roping = 0. Roping is deemed to be occurring if one or more of the roping criteria as stated in Equation 2-1, Equation 2-2, and Equation 2-3 are met. Figure 5-1 illustrates the simulated open

circuit hydrocyclone system with the status of variables available to the RL controller noted.

Figure 5-1

Simulated Open Circuit Hydrocyclone System Diagram



Note. Simulated open circuit hydrocyclone system illustration showing the relevant variables, along with their visibility and nature in relation to the RL controller.

5.1.2 Reinforcement Learning Controller

Lau's (2016) implementation of DDPG (Lillicrap et al., 2015) was used with permission as a starting point for the RL control. One point of difference to note is that Lau's (2016) implementation does not contain batch normalisation in the neural networks, unlike Lillicrap et al.'s (2015) DDPG architecture. Lau's (2016) implementation uses Keras (Chollet & others, 2015a) with a Tensorflow (Abadi et al., 2015) back end, and was originally interacting with a car driving simulator. Both Lau's (2016) and Lillicrap et al.'s (2015) implementations of DDPG use the optimiser Adam (Kingma & Ba, 2014) and ReLU activation for the hidden fully connected layers of both actor and critic networks. Additionally, the hyperparameters below are consistent with both Lau (2016) and Lillicrap et al. (2015) implementations of DDPG

- Discount factor $\gamma = 0.99$
- Target networks' soft update $\tau = 1 \times 10^{-3}$

- Actor Learning Rate = 1×10^{-4}
- Critic Learning Rate = 1×10^{-3}

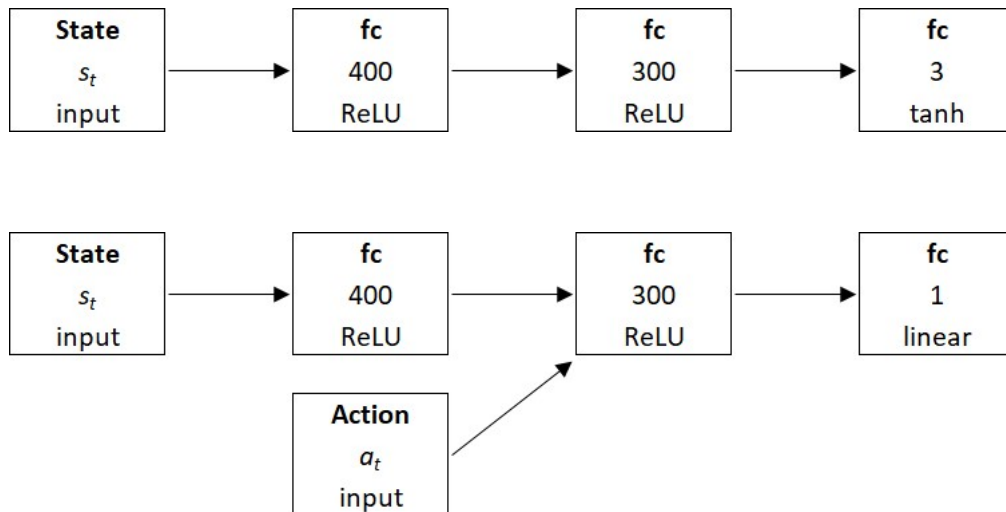
Lau's (2016) DDPG implementation was modified as required for this investigation including

- For use with Tensorflow's own implementation of Keras.
- The actor network was brought closer in line with Lillicrap et al.'s (2015) DDPG architecture, with the state input connected to the first 400 node fully connected layer followed by a second 300 node fully connected layer. The output of the network was then a three node fully connected layer with tanh activation for each node to produce the action output. As illustrated in Figure 5-2.
- The critic network was brought closer in line with Lillicrap et al.'s (2015) DDPG architecture, with the state input connected to the first 400 node fully connected layer and the actions not introduced until the second 300 node fully connected layer. The output of the network was then a one node fully connected layer with linear activation for the critic network to output the $Q(s_t, a_t)$ value. As illustrated in Figure 5-2.
- While Lillicrap et al.'s (2015) DDPG implementation used 1×10^{-2} L2 weight decay for the critic, the implementation noted in Fujimoto, van Hoof, and Meger (2018) did not. For this implementation 1×10^{-2} L2 kernel regularisation was used on the critic output layer.
- The actor and critic network's kernel initializer for their respective final layers were chosen from a random uniform distribution -3×10^{-3} – 3×10^{-3} as in Lillicrap et al. (2015).
- Following Fujimoto, van Hoof, and Meger (2018) demonstration of reducing per-update error the actor policy and target networks' updates were delayed, occurring every second time step while the critic continues to be updated every time step.
- The replay buffer size was selected to be large enough to hold the agent's entire history for a given test scenario, as was used in Fujimoto, van Hoof, and Meger (2018).

- A mini-batch size of 100 was selected as per Fujimoto, van Hoof, and Meger (2018).

Figure 5-2

Actor and Critic Neural Network Architectures



Note. Actor (top) and critic (bottom) neural network architectures used in the RL controller, incorporating aspects of Lillicrap et al. (2015) and Lau (2016) implementations of DDPG. Includes the number of nodes and activation for fully connected (fc) layers.

5.1.3 Simulator-Control Interaction

At the start of each episode the hydrocyclone feed flow rate, hydrocyclone feed solids concentration, tank level and apex diameter were selected from a random uniform distribution within the initialisation range. The standard initialisation range shown below was used unless otherwise stated.

- $Q = 1000\text{--}2000 \text{ m}^3/\text{h}$
- $C = 40\%\text{--}60\% \text{ w/w}$
- $T_L = 40\%\text{--}60\%$
- $D_A = 90\text{--}140 \text{ mm}$

These selected values were then used to set the tank solids feed rate for the episode, under the simplified assumptions that inflow volume equals outflow volume at the start of an episode and a constant feed solids density (unless otherwise stated). The SPs and feed sizing %passing modification value were then set for the episode, depending on the experimental requirements.

The input values $\{ T_{L[\text{initial}]}, Q, C, D_A, R_f, \rho_s, M \}$ were then fed to the simulator and the output values $\{ T_{L[\text{final}]}, Q, C, D_A, P, X, U_P \}$ returned. The output values for the simulator consist of variables typically measured in a hydrocyclone circuit. One exception to this being hydrocyclone operating state detection, though image-based (Giglia & Aldrich, 2020) or vibration-based (Bowers et al., 2019) sensors are potential options. The other exception being hydrocyclone underflow particle sizing information, in this case hydrocyclone underflow P80. Though vibration-based sensors (Cahill, 2021) are available if physical online measurements of the composite stream was not occurring. An image-based sensor like that explored in Section 4 could also be another option if successfully trained.

The underlying values interacting with the simulator were maintained separately from the observations visible to the controller. This allows for the injection of noise (to simulate sensor noise) to be applied. The state representation is then formed from these noisy observations of the simulator output and the SPs, resulting in the form $(T_L, Q, C, D_A, P, X, U_P, T_L - T_{L[\text{SP}]}, U_P - U_{P[\text{SP}]})$. This representation of state using observations and SP difference is similar to that presented by Hafner and Riedmiller (2011). Observations of the simulated system continue to be received, and along with the SP differences, stacked over time in the state representation with no control actions taken until the set state matrix size has been filled. A step will thus refer to an action taking time step after the state matrix is full.

Gupta and Eren's (1990) laboratory hydrocyclone control experiments discussed in Section 2.6.2 were used as inspiration for this experimental design. In their case, adjustments could be made to the following variables sequentially to restore the cut size to a target value

- Vortex finder height, via a uniquely designed pneumatic system

- Apex diameter, via a pneumatically adjustable rubber sleeve
- Hydrocyclone feed flow rate, via a variable speed pump

For the simulated control experiments three continuous actions were selected to be able to be adjusted simultaneously at each time step

- Hydrocyclone feed flow rate (Q), which in practice would be adjusted via pump speed
- Hydrocyclone feed solids concentration (C), which in practice would be adjusted by dilution water addition
- Apex diameter (D_A), which could be made adjustable via a pneumatically controlled rubber sleeve as noted by Gupta and Eren (1990)

The flattened current state matrix values were rescaled by dividing each variable by a constant as shown $\left(\frac{T_L}{100}, \frac{Q}{3000}, \frac{C}{100}, \frac{D_A}{200}, \frac{P}{200}, \frac{X}{1}, \frac{U_P}{3000}, \frac{T_L - T_{L[SP]}}{100}, \frac{U_P - U_{P[SP]}}{3000}\right)$, before being used as input to the actor network. Note that all state values were rescaled in this way prior to being used in the actor and critic networks. The actor network then outputs the three action values, bounded -1.0–1.0 due to their tanh activation functions. It is worth noting that rather than using squashing functions (e.g., tanh or sigmoid) to bound the actions, Hausknecht and Stone (2015) used an inverting gradient method. In their method gradients were modified as parameter boundaries were approached and was found to be necessary for stable learning in their domain, avoiding issues associated with squashing functions saturating. Spielberg (2017) also used the inverting gradient method for their process control implementation, though in the current work stable learning was possible using tanh activation (with clipping if action exploration suggests a value outside these bounds as will be discussed later) and thus modifying gradients was avoided.

During training, exploration noise is then added to each action. Fujimoto, van Hoof, and Meger (2018) added Gaussian exploration noise (mean = 0, SD = 0.1) having noted no performance benefit in using temporally correlated noise as in the Lillicrap et al.'s (2015) DDPG implementation. For this investigation, a more aggressive Gaussian exploration noise (mean = 0, SD = 0.2) was applied during training. The randomly initialised policy was used for the first 2,000 time steps (with Gaussian

noise added) before training of the networks commenced, thus initialising the replay buffer with a broad set of transitions to begin training on.

As the actor and critic networks were randomly initialised, their early outputs during training were not yet specialised for the system. To reduce the likelihood of exceeding what would be physical limits of an actual circuit, particularly during training, safety limits were selected. The current MV and proposed actions' effects were assessed against the safety limits below.

- $Q = 500\text{--}2500 \text{ m}^3/\text{h}$
- $C = 30\%\text{--}70\% \text{ w/w}$
- $D_A = 70\text{--}160 \text{ mm}$

Given the current value, if a proposed action (plus exploration noise if applied) was determined to exceed a safety limit, then the action was modified so that its calculated subsequent value is on the safety limit. In this way for a sensor noise free environment the control would not exceed these safety limits, though in a noisy environment it was still able to slightly exceed the limit but will be forced to try to return to the limit at each subsequent time step. The final stage of action interrogation was that if the proposed action value was > 1.0 or < -1.0 , then the values 0.9999 or -0.9999 were forced, respectively. This stopped the action boundary being exceeded (which was only possible during training due to the addition of exploration noise) and by using fixed forced values this aided identifying when it had been applied.

In Syafiie et al.'s (2008) Q-learning based control with discretised states and actions for chemical processes, the actions available in each state were limited in part by the allowable incremental variation in the MV. The MV would be adjusted by multiplying the gain by the desired action and adding it to the previous MV value. In a comparable manner for this investigation, the action values are then multiplied by the relevant maximum step size, being the maximum allowable change in MV per time step given the continuous action values are bound between $-1.0\text{--}1.0$. This is then added to the relevant MV's value, and these revised inputs fed to the simulator. The subsequent noisy observations and SP differences were then stacked over time

in the state representation, with the oldest value removed to maintain the set size for the subsequent state representation.

The transition reward was then calculated based on the subsequent state. A simple form of reward function for process control solely gives a penalty when the system is outside of the SP tolerance range (Hafner & Riedmiller, 2011). Though by shaping the reward function (Ng et al., 1999) richer feedback can be provided to aid training. This can consider difference from SPs (Spielberg, 2017), penalties for undesirable conditions (i.e., exceeding safety limits or roping state), and the relative weightings of rewards. The reward structure used was as follows

- - $|T_L - T_{L[SP]}| / 200$
- - $|U_P - U_{P[SP]}| / 5$
- - 5 if roping ($X = 0$)
- - 0.1 if equal to or exceeding safety limits

These values were selected so that in typical operating ranges the total reward values are not excessively large. Table 5-1 illustrates the reward corresponding to a range of underflow P80 and tank level SP difference conditions. The tolerance values for underflow P80 and tank level used in the experiments are shown in bold in Table 5-1. The reward at tolerance for tank level results in a greater penalty than at tolerance for underflow P80, the intent being to give a stronger weighting to flow stabilisation than product size, though the training process looks to maximise total return thus it is not necessarily a direct trade off. The penalty for roping has been set beyond what would be expected for typical tank level and underflow P80 SP differences, to strongly discourage the policy from approaching regions where roping could occur. There was also a small penalty when the MV's safety limits were met or exceeded, this is intended to encourage the policy itself to operate near these limits without the hard coded safety intervention from taking place.

Table 5-1

Example Process Variable Set Point Differences and Corresponding Reward

Reward	$ U_P - U_{P[SP]} $ (μm)	$ T_L - T_{L[SP]} $ (%)
-0.10	20	0.50
-0.25	50	1.25
-0.50	100	2.50
-1.00	200	5.00
-2.00	400	10.00
-5.00	1000	25.00

Note. Resulting reward from absolute difference tank level reading (T_L) to relevant set point ($T_{L[SP]}$), and underflow P80 (U_P) to relevant set point ($U_{P[SP]}$).

The flattened state matrix values, action values, reward, and flattened subsequent state matrix values are then stored in the replay buffer. Critic, actor, and target network updates were then performed as required. The control process then continued until a terminal condition was reached for the episode. If the tank level was deemed to have lost control (i.e., subsequent noisy observation <5% or >95%) then the current episode would terminate. Otherwise, the terminal conditions were either that for last 10 steps the current sensor readings had been within the tolerance range of the SPs and roping had not occurred, or 50 steps had occurred, whichever came first.

5.2 Case studies

5.2.1 State Representation

RL control is ideally structured around an MDP, where the state is fully observable. While in a real-world setting sensor information is not typically available to cover all relevant contributing variables, so the system is partially observable. For Mnih et al.'s (2015) DQN implementation partial observability was handled by including a history of observations and actions over a fixed time horizon (x) in the state representation, $s_t = [o_{t-x}, a_{t-x}, \dots, a_{t-1}, o_t]$. Similarly, for Lillicrap et al.'s (2015) DDPG implementation when learning direct from the rendered environment, one

controller time step constitutes three environment renderings with the action taken repeated during this period. The collection of feature maps from the renders are then used in the state representation, where changes in frames allows information like velocity to be inferred.

To trial the control behaviour of a simple industrially applicable partially observed system, the tank solids feed rate was randomly set and unmeasured at the start of each training episode. All other variables were either measured or held constant. A state representation was formed by stacking observations (sensor readings) along with difference between current value and SP over a specified time horizon. The historic actions taken do not explicitly form part of the state representation, but as the variables they are applied to were measured the changes are still implied in the state representation.

Model performance after a given number of training episodes, with the state representation made up of only the observations and SP differences for the current time step (State 1) and the observations and SP differences for the current time step along with the previous two time steps (State 3), are shown in Table 5-2. Testing then occurred on 1,000 episodes, with random number seeding meaning the initial conditions for each episode was consistent for all test scenarios. The action taking steps were after the state representation was full, so for State 1 this was after the first sensor reading. For State 3 no actions were taken until the state matrix was full, though given for these tests the initialisation tank inflow = outflow it was effectively three sets of identical observations and SP differences before action taking steps occur. The training and testing SPs were $T_{L[SP]} = 50\%$ and $U_{P[SP]} = 1800 \mu\text{m}$.

Table 5-2*State Representation and Training Period Test Performance Metrics*

State	Training Episodes	Episodes Solved (<50 steps)	Average Steps (for solved episodes)
State 1	1,000	75.6%	21
	1,500	71.2%	21
	2,000	49.1%	23
	2,500	74.7%	22
State 3	1,000	100.0%	18
	1,500	100.0%	18
	2,000	100.0%	19
	2,500	100.0%	18

Note. Percentage of episodes terminated in <50 steps and the average steps taken to solve them, for different state representations and periods of training.

The results in Table 5-2 show that a state representation of three consecutive time steps was sufficient to allow the model to infer missing information and solve the test episodes in <50 steps after only 1,000 episodes worth of training. Thus, the three consecutive time steps state representation (State 3) will continue to be used for subsequent experiments. Also, training for 2,000 episodes worth of experiences was found to provide an acceptable level of testing performance over a range of test conditions encountered in the subsequent experiments. Thus, this was used as the testing point for all trained models unless otherwise stated, having considered the potential for issues such as overfitting and catastrophic forgetting (Goodfellow et al., 2013).

5.2.2 *Unsolvable Instances*

The training process was repeated under the same conditions as Section 5.2.1 (including State 3 representation), but with a 2000 μm hydrocyclone underflow P80 SP. Figure 5-3 shows the number of steps taken for the test episodes to terminate, with 89.7% of episodes terminating in <50 steps. Figure 5-4 also shows the initial and final (at episode termination) tank level and hydrocyclone underflow P80 SP difference (PV – SP) values for the 1,000 test episodes. By looking into the 103

episodes that did not terminate in <50 steps we find that the final tank level reading ranged from 50.1%–50.3%, thus within the tolerance range. Though the final hydrocyclone underflow P80 reading ranged from 1817.2–1950.1 μm , thus either outside of the tolerance range or not continuously within the tolerance range for sufficient time to allow the episode to terminate in <50 steps.

Figure 5-3

Count of Test Episodes Terminated in the Given Number of Steps





Figure 5-4

Initial and Final Process Variable Set Point Differences for Unsolvable Instances Test



Note. Set point difference is process variable subtract set point ($PV - SP$).

With access to the underlying simulator's roping criteria the ratio of apex diameter to vortex finder diameter that leads to roping is known (Equation 2-1). Given the fixed vortex finder diameter of 190.5 mm, then an apex diameter less than 85.725 mm would lead to roping. The final apex diameter of the episodes that did not terminate in <50 steps ranged between 86.9–87.3 mm. For all these episodes the tank solids feed rates were low (Figure 5-5), with subsequent lower hydrocyclone feed flow rates and feed solids concentrations reached to balance the circuit volumetric flow. Under these conditions the control can only reduce the apex diameter so far, to increase the hydrocyclone underflow P80, whilst avoiding the potential of roping. Figure 5-6 shows the change in tank level and hydrocyclone underflow P80 over time for the lowest tank solids feed rate test episode (#82, $R_f = 545.5$ dry t/h) and the highest tank solids feed rate test episode (#792, $R_f = 787.3$ dry t/h) which failed to solve in <50 steps.

Figure 5-5

Test Episode Tank Solids Feed Rate Against Steps Taken for Episode to Terminate





Figure 5-6

Process Variables Plot Over Time for the Nominated Test Episodes



Note. Change in tank level (T_L) and underflow P80 (U_p) over time (t) for test episodes #82 and #792 which failed to solve in <50 steps.

Of note, for the 24,450 action taking transition made during this test, none resulted in a roping state. Taken together the behaviour presented for this simple partially observed system is quite desirable for process control. The policy generates actions to approach the target conditions, while if the target conditions are not achievable then operates near the limit and ideally does not enter any undesirable operating regions.

5.2.3 Partial Observability Extension

The Section 5.2.1 training process with State 3 representation was again used, with feed particle size included as an addition unmeasured variable randomly set at the start of each episode. To facilitate this a random value between -0.8–0.8 was selected and added to each screen size's percentage passing value of the base case hydrocyclone feed particle size distribution, with this value referred to as feed sizing

%passing modification (M). In this way the feed particle size is shifted, resulting in a range of possible feed size distributions with the extremes shown in Table 5-3.

Table 5-3

Summary of Hydrocyclone Feed Sizing Information for Modified Base Case

Feed Sizing	Base – 0.8 (μm)	Base (μm)	Base + 0.8 (μm)
P80	1368	1273	1185
P50	290	281	273

Note. Base case from Sepúlveda (2012) $\pm M$.

For these experiments three random number seeds were used during training (seed = 10, 20, 30), and the resulting policies after 2,000 episodes of training are compared on the same random number seeded 1,000 test episodes. The test performances appear similar (see Table 5-4, Figure 5-7, and Figure 5-8), with minor differences as would be expected considering the training is performed on episodes with different initial conditions, transitions, and ultimately number of gradient updates. All policies were unable to solve the same seven episodes in <50 steps, with training seeds 20 and 30 also unable to solve the same additional three episodes. The initial conditions for these test episodes were all $R_f < 675$ t/h and $M > 0.25$. This presents a similar issue as covered in Section 5.2.2 where at lower throughputs, and in this case also coupled with finer feed sizing, the control is limited in how close to conditions that will lead to roping it will approach (final apex diameters were between 85.8 mm to 88.5 mm).

Also of note is that not all trained policies were able to result in no roping events. All these events occurred due to apex diameter adjustment resulting in a value of $D_A < 85.725$ mm, highlighting that these policies did not behave optimal under all conditions. Though considering the approximately 19,000 transitions taken during each 1,000 episode test, this level of imperfect behaviour would likely be acceptable for this circuit, though not appropriate for a safety critical system.

Table 5-4


Test Performance Metrics for Policy Trained With Different Random Number Seeds

Training Seed	Episodes Solved (<50 steps)	Average Steps (for solved episodes)	Roping Events
10	99.2%	19	8
20	98.9%	18	0
30	98.9%	19	9

Note. Percentage of episodes terminated in <50 steps and the average steps taken to solve them, along with number of roping events, for different training random number seeds.

Figure 5-7


Count of Test Episodes Terminating in the Given Number of Steps



Note. 1,000 test episodes for each training random number seed.

Figure 5-8

Initial and Final Process Variable Set Point Differences for Training Seed Tests



Note. Set point difference is process variable subtract set point ($PV - SP$).

Another 1,000 episode test using the trained model (seed = 10) in conditions outside that experienced during training was undertaken to test the generalisation performance. For this test, the initial tank level was randomly selected between 30%–40% or 60%–70% (outside of the training initialisation range 40%–60%). The feed sizing modification for each episode was randomly selected between -1.0 to -0.8 or 0.8 to 1.0 (outside of the training range -0.8–0.8), see Table 5-5 for the extreme feed sizing information. The tank inflow volume equals outflow volume at the start of an episode was also not applied during these tests. Instead, the tank solids feed rate for the episode was randomly set between -200–200 t/h different to the solids discharge rate out of the tank (which is based on the randomly initialised hydrocyclone feed flow rate and solids concentration). Though the tank solids feed rate was clipped between 500–2000 t/h, just outside the approximate range of -538.5–1953.5 t/h encountered during training.

Table 5-5

Summary of Hydrocyclone Feed Sizing Information for Extreme Modified Base Case


Feed Sizing	Base – 1.0 (µm)	Base + 1.0 (µm)
P80	1392	1160
P50	292	270

Note. Base case from Sepúlveda (2012) $\pm M$.

Testing resulted in 96.2% of the episodes solved in <50 steps, with initial and final SP difference shown in Figure 5-9. The remaining 38 episodes were all in the lower throughput coupled with finer feed sizing region for which reaching target underflow P80 of 1800 µm may not be possible without roping. All the final time step's tank levels for these unsolved episodes were either within or near the tolerance region (specifically 47.4%–49.3%). In this case the region where $R_f < 705$ t/h and $M > 0.8$ contained the unsolvable episodes (Figure 5-10). The higher tank solids feed rates and all coarser feed sizing episodes were solvable as they did not require approaching the apex diameter that leads to roping. Of the 22,630 steps taken in these test episodes 283 resulted in roping events, all involved $M > 0.8$ and were the result of $D_A < 85.725$ mm roping threshold. Of the roping events 262 occurred with $R_f < 705$ t/h, making this level of roping occurrences not unreasonable given that this is an area of state space not explored during training and with unmeasured variables known to cause challenging control conditions.

Figure 5-9


Initial and Final Process Variable Set Point Differences for Generalisation Tests



Note. Set point difference is process variable subtract set point ($PV - SP$).

Figure 5-10

Solvability of Test Episodes Considering Feed Size Modification and Solids Into Tank



Note. Solved refers to episode terminated in <50 steps. Only the finer feed size modification region is shown, as all coarser feed size modifications tested were solved in <50 steps.

5.2.4 Set Point Variability

The controller so far has been trained using fixed SPs $T_{L[SP]} = 50\%$ and $U_{P[SP]} = 1800 \mu\text{m}$. To see how the trained system from Section 5.2.3 (seed = 10) handles variable SPs a 1,000 episode test was performed with the tank level SP varied between 40%–60% (inclusive) and the hydrocyclone underflow P80 varied between 1700–1900 μm (inclusive). For the Coarse test case the tank level SP was randomly selected in 5% increments, and the hydrocyclone underflow P80 SP randomly selected in 50 μm increments. For the Fine test case the tank level SP was randomly selected in 1% increments, and the hydrocyclone underflow P80 SP randomly selected in 10 μm increments.

Table 5-6 presents key performance metrics for the variable SP test from the fixed SP trained model. The number of episodes able to be solved in <50 steps with variable SPs is significantly less than that achieved with the fixed SPs (for which it was trained). While measured tank level and hydrocyclone underflow P80 value along with the SP difference values are used to form the state representation, because in the fixed SP case the SP differences values are always just a constant offset of the measured values the model does not learn the effect of different SP scenarios.

Table 5-6

Test Performance Metrics for Policy Trained With Fixed Set Points

SP Variability	Episodes Solved (<50 steps)	Average Steps (for solved episodes)	Roping Events
Fixed	99.2%	19	8
Coarse	58.4%	20	102
Fine	65.8%	19	28

Note. Percentage of episodes terminated in <50 steps and the average steps taken to solve them, along with number of roping events, for policy trained with fixed set points (SPs) and tested with different SP variability.

The training process was repeated for 2,000 episodes, but this time the Coarse range of SPs were randomly set at the start of each episode and held constant. Testing of the trained policy for 1,000 episodes with varying SP scenarios was then performed, with key metrics summarised in Table 5-7. This trained model was able to handle the Coarse SP variability and extends well to the Fine SP variability scenario. It also had only slightly lower performance on the fixed SP test case than the model developed solely on this scenario. This demonstrates that if target SPs can or will be varied then this should be incorporated into the training process, though it does add extra complexity.

Table 5-7

Test Performance Metrics for Policy Trained With Varied Set Points

SP Variability	Episodes Solved (<50 steps)	Average Steps (for solved episodes)	Roping Events
Fixed	97.6%	19	5
Coarse	95.8%	20	20
Fine	95.5%	20	1

Note. Percentage of episodes terminated in <50 steps and the average steps taken to solve them, along with number of roping events, for policy trained with Course set point (SP) variability and tested with different SP variability.


In the tests so far, the SPs were selected at the start of each episode and held constant, with the difference included for each time step making up the state representation. If SP variability was allowed on an ongoing basis, as would occur in continuous process control, the state representation needs to accommodate this. Hafner and Riedmiller (2011) noted that a SP change violates the Markov property of the MDP model, but by using the current SP at each time step (t) and using that same value in forming the subsequent time step's ($t+1$) state representation, they propose a consistent and valid transition can be built. Considering this requirement, the current state representation with a three time step history, and constant SPs during training, then the varying SP condition was considered as follows.

The SPs (SP_t) used in the current state (s_t) representation (thus compared against $[o_{t-2}, o_{t-1}, o_t]$) are still applied in determining the reward (r_t) and the subsequent state (s_{t+1}) representation (thus compared against $[o_{t-1}, o_t, o_{t+1}]$) to complete the (s_t, a_t, r_t, s_{t+1}) transition. The historic observations $[o_{t-1}, o_t, o_{t+1}]$ are then reassessed against the new SPs (SP_{t+1}) prior to forming the new current state for the system. The model trained with varied SPs was then used in a test episode in which the SPs were changed at certain intervals. All other random variables selected at the start of the episode remained constant. Figure 5-11 shows the change in the PVs tank level and hydrocyclone underflow P80, along with the respective SPs, during the episode. Figure 5-12 shows the change in the MVs hydrocyclone feed flow rate,

hydrocyclone feed solids concentration and apex diameter during the episode. The episode terminated successfully after holding the system within the SP tolerance range.

Figure 5-11


Process Variables and Relevant Set Points During the Set Point Change Test



Note. Change in tank level (T_L) and underflow P80 (U_p) over time (t) with varied set points (SP).

Figure 5-12

Manipulated Variables During the Set Point Change Test



Note. Change in hydrocyclone feed flow rate (Q), hydrocyclone feed solids concentration (C) and apex diameter (D_A) over time (t) for the set point change test.

If offline pretraining was to be performed on historic transitions, SP variability could be incorporated. With the collection of sensor observations over the required time period, the actions taken, and the subsequent sensor observations then (s_t, a_t, r_t, s_{t+1}) transitions for varying SPs can be formed. Reasonable SPs for the system can be selected and used in forming the state representations. The subsequent state can then be assessed against the defined reward function to give the transition reward. In this way, regardless of the underlying control strategy and actual SPs that formed the observations, these synthetic transitions could potentially be used in the off-policy training method.

5.2.5 Noise

All experiments to this point have the measured variables true to the underlying system, however uncertainty in sensors readings would occur in an

industrial setting. To simulate sensor noise random values were added to the simulator output values prior to forming the state representation. The random values were drawn from a normal distribution with mean = 0 and SD as shown in Table 5-8, with the Noise B having triple the SD of Noise A.

Table 5-8


Standard Deviation of Random Noise Applied to Measured Variables for Each Test

Variable	Noise A	Noise B	Units
Tank level	0.2	0.6	%
Hydrocyclone feed flow rate	5	15	m ³ /h
Hydrocyclone feed solids concentration	0.15	0.45	% w/w
Apex diameter	0.3	0.9	mm
Inlet pressure	0.3	0.9	kPa
Operating state (0 or 1)	0	0	
Underflow P80	5	15	μm

Both training and subsequent testing of the trained model were performed with the relevant sensor noise condition applied. The initial training and testing was performed using the settings as in Section 5.2.3, with fixed SPs and three observations used in the state representation (State 3). The initial and final SP difference plots for the Noise A and Noise B 1,000 episode tests are shown in Figure 5-13 and Figure 5-14 respectively. Table 5-9 shows the key performance metrics for these tests, along with performance metrics for the trained model from Section 5.2.3 (training seed = 10) on the noise free version of this test system. With increasing extent of noise, the number of episodes able to be solved in <50 steps decreases and the average steps taken for solved episodes increases.

Figure 5-13


Initial and Final Process Variable Set Point Differences for State 3 Noise A Tests



Note. Set point difference is process variable subtract set point (PV – SP).

Figure 5-14

Initial and Final Process Variable Set Point Differences for State 3 Noise B Tests



Note. Set point difference is process variable subtract set point (PV – SP).

Table 5-9

Test Performance Metrics for Policies Trained With Varied Noise Conditions

State	Noise Condition	Episodes Solved (<50 steps)	Average Steps (for solved episodes)	Roping Events
State-3	No noise	99.6%	19	8
	Noise A	83.3%	20	3
	Noise B	9.3%	32	8
State-5	Noise A	95.7%	21	10
	Noise B	87.7%	27	29


Note. Percentage of episodes terminated in <50 steps and the average steps taken to solve them, along with number of roping events, for different state representations and noise conditions.

The noise introduces uncertainty into the state representation and degrades the training process and ultimate model performance. By considering noise as further partial observability, a simple extended state representation using the last five observations (State 5) was trialled to see if this could provide further information to improve control. Table 5-9 also shows the key performance metrics for the State 5 tests. This demonstrated an improved number of episodes solved and lower average steps for solved episodes compared to the State 3 noise tests, with a slight increase in roping events occurring.

Along with the presence of noise brings uncertainty into defining the system's state, there is also a requirement to be operating closer to the SPs to ensure even just the sensor noise alone does not result in a reading exceeding the tolerance range for the 10 steps required for the episode to terminate. Considering this, though performance is not as strong as the noise free system these results are still reasonable. Figure 5-15 shows the feed sizing %passing modification and tank solids feed rate conditions under which the State 5 test episodes were not solved in <50 steps. For Noise A it was mainly the finer feed sizes coupled with the more extreme low or high tank solids feed rates that were not solved. While for the Noise B test, with a higher degree of applied noise, a wider range of conditions were not solved consistent with the increased difficulty in underlying system state identification and tolerance requirements.

Figure 5-15

Unmeasured Variables for Unsolved State 5 Variable Noise Test Episodes




Note. Solved refers to episode terminated in <50 steps.

The initial and final SP difference plots for the Noise A and Noise B 1,000 episode tests for the State 5 representation are shown in Figure 5-16 and Figure 5-17 respectively. Most final SP difference results are within or near the tolerance limits, however two points in the Noise B test resulted in final tank level difference values of -9.5% and -7.5% (episodes #106 and #800 respectively) and will be examined further.

Figure 5-16


Initial and Final Process Variable Set Point Differences for State 5 Noise A Tests



Note. Set point difference is process variable subtract set point (PV – SP).

Figure 5-17

Initial and Final Process Variable Set Point Differences for State 5 Noise B Tests




Note. Set point difference is process variable subtract set point ($PV - SP$).

Episode #106 had $R_f = 551.1$ t/h and $M = 0.76$, while episode #800 had $R_f = 570.0$ t/h and $M = 0.61$. Both episodes are near the low limit of $R_f = 538.5$ t/h experienced during training and coupled with a fine feed size, making $U_{P[SP]} = 1800$ μm unattainable due to the roping apex diameter limit, as discussed in Section 5.2.3. Figure 5-18 shows the change in the PVs tank level and hydrocyclone underflow P80 for these episodes. Figure 5-19 shows the change in the MVs hydrocyclone feed flow rate, hydrocyclone feed solids concentration and apex diameter for these episodes.

Figure 5-18


Process Variables Plot for State 5 Noise B Tests With Decreasing Tank Level



Note. Change in tank level (T_L) and underflow P80 (U_p) over time (t) for test episodes #106 and #800 with decreasing tank level.

Figure 5-19

Manipulated Variables Plot for State 5 Noise B Tests with Decreasing Tank Level



Note. Change in hydrocyclone feed flow rate (Q), hydrocyclone feed solids concentration (C) and apex diameter (D_A) over time (t) for test episodes #106 and #800 with decreasing tank level.

For both tests, the controller rapidly reduces the apex diameter to near, and occasionally below, the roping limit. The hydrocyclone feed flow rate is also rapidly reduced to around the 500 m³/h minimum flow limit. Both behaviours are expected given the fine feed size and low tank solids feed rate. The hydrocyclone feed solids concentrations are also rapidly increased to the 70% w/w upper limit; however, this is too high and leads to the tank level dropping. A cycling of the hydrocyclone feed solids concentration between around 60% w/w to the 70% w/w limit then begins, which still leads to the tank level dropping. Ultimately, near the low limit of feed rate, in an unsolvable region given the feed particle size distribution, and with this level of noise there were thus conditions encountered through which the current trained controller had not reach a stable policy for this time frame. This highlights that it would still be prudent to have control override protections in place for atypical or

undesirable conditions, such as water addition to ensure the tank level does not drop below a lower safety limit.

5.3 Discussion

The following section will highlight further limitations and points of interest in applying RL techniques to mineral processing circuit control, along with works related to these areas.

5.3.1 System

For the simulated environment the controller interacted with, the effect of actions was instantaneous and held constant for the entire time step. The unmeasured variables were also constant for an entire episode, thus constant within each time step. While this was applied for simulation simplicity, in practice the time step for high level control would need to be selected considering the process dynamics so that it is sufficient for the actions and their effect to be realised, whilst also frequent enough to be able to compensate for changing conditions.

Tallec et al. (2019) found that off-policy RL methods using the action-value function (with DQN and DDPG tested) were affected by the time discretization used, and subsequently extent and frequency of actions taken. At the extreme, they noted that when approaching continuous-time the strength of a briefly maintained action over another became increasingly small leading to poor Q-learning performance. As an industrial example, for Guo et al.'s (2019) grinding circuit RL control investigation the available historic data was processed to determine the most appropriate transition time step to use. The sensor readings, originally sampled at 5 s intervals, had 1 min moving averaging applied. Correlation analysis was then performed to determine the optimal time interval for the control action to influence the system.

For the simple open circuit system in these experiments, the stacking of a few consecutive sensor readings and SP differences was able to provide largely sufficient information for control over the range of conditions tested. Though improved and more widely applicable representations could be developed. For more complex systems with longer time delays or interactions, such as a closed circuit grinding circuit where the hydrocyclone underflow is returned to the mill along with the fresh

feed, further adaptations to the RL controller design would likely be required. The use of recurrent neural networks (RNNs) such as Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997) for handling sequential data may be beneficial. RNNs incorporate a hidden state to represent recent input data, and LSTM uses gates which learn to control the flow of relevant data along with a cell state designed to stabilises error flow during back propagation training (Hochreiter & Schmidhuber, 1997; Wang et al., 2017). Wang et al. (2017) demonstrates the potential for incorporating LSTM into both networks of an on-policy actor-critic RL controller for a simulated buildings' heating ventilation and air conditioning system. They also noted the potential benefit LSTM's have in handling partial observability and noise through the hidden state's inclusion of historic data to aid in true state inference.

As we move towards more realistic processing systems the associated complexity, partial observability, and noise make the attainment of a state representation that satisfies the Markov property increasingly difficult. With then only an approximation of state the longer-term predictive performance can suffer, though if we accept this relaxed approximation of state good/sufficient predictions may still be attainable (Sutton & Barto, 2018). The performance of the State-5 Noise-B tests in Section 5.2.5 highlights this point, given that largely successful control was still attained though long-term decreasing tank level events occurred under certain conditions. Considering the detrimental effect that increasing levels of noise yielded, the application of reasonable filters on the sensor readings may also be beneficial. While every effort should be made to characterise the system as well as practicable, RL based control methods need to be able to adequately handle this approximate state if they are to be successfully applied to live processing plant control.

5.3.2 Actions

For the experiments undertaken three continuous actions were able to be adjusted at each time step. As discussed in Section 5.1.3, these were applied with a maximum step size which restricted how much the MVs can change from their previous values at each time step. This was a particularly important consideration for the mineral processing industry given the requirement for stable slurry transport.

Following the selection of a suitable time step for the controller, the maximum step size for each action could then be selected considering knowledge of the process and an analysis of historic transitions to determine what would be considered reasonable and safe.

For systems where only one action is desired or permitted per time step, the control problem becomes that of action selection followed by corresponding continuous parameter selection. As noted in Section 2.7.2, Hausknecht and Stone (2015) presented an extension of DDPG for these parametrised action-space systems. Their actor network used separate output layers, one having a node for each action and the other having a node for each action's associated parameter. The applied action was then the maximum value from the action layer and its associated parameter/s from the parameter layer. While all values of the actor output layer are provided to the critic, and thus the applied action is not explicitly indicated, their success in training demonstrates that this was still a viable method.

Another common type of system that would be encountered in mineral processing plant control would be where both continuous and discrete control actions are available. Hallén et al.'s (2019) RL grinding circuit control experiments included the potential for a discrete action to divert a flap gate, along with the other continuous control actions. Their RL controller did not manage to attain a stable control strategy for the flap gate, while it was utilised in the existing PID control strategy used for performance comparison. In the case of hydrocyclone circuits, the switching of hydrocyclones in and out of operation would be a key discrete action to have available to the controller. Though along with this comes another layer of complexity as the controller must learn the behaviour of the circuit under a range of conditions and relevant number of operating hydrocyclones. Then it must learn when to bring a hydrocyclone in/out of operation for best long term expected return and allow sufficient time for the system to stabilise following such an abrupt change (to avoid cycling of hydrocyclone operation).

5.3.3 Training

In the experiments performed, training was through interaction with a simulated environment, starting from randomly initialised networks, and allowed to explore the environment widely. To train RL controllers for a live mineral processing plant, it is unlikely that such a degree of random initial behaviour and exploration range of positive and negative scenarios would be acceptable given the potential cost of lost production and risk to equipment and personnel that may arise. At a minimum, a degree of offline pretraining of the control would need to be performed. Should a suitable model of the system be available then pretraining could be performed against this, but if a sufficiently developed model is already available then it would need to be determined if model-based control methods may be better to serve the system.

A model free off-policy method that incorporates a replay buffer, such as DDPG (Lillicrap et al., 2015), does lend itself to be suitable for offline pretraining from a replay buffer constructed from historic operating data. While the naïve approach would be to simply train the networks directly off such a replay buffer of historic transitions, the results from Fujimoto, Meger, and Precup (2018) suggest this may not lead to ideal behaviour. They performed a series of experiments involving a secondary DDPG controller learning in a true off-policy manner from

- the complete final replay buffer used in training a primary DDPG controller, with high exploration noise
- the replay buffer of a primary DDPG controller as it is training on the system, thus the same evolving dataset is used to train both controllers
- a replay buffer built from transitions produced by an already trained DDPG controller

Fujimoto, Meger, and Precup (2018) found all of these scenarios resulted in the trained secondary controller demonstrating worse performance than the alternative controller. They attributed the performance deterioration to extrapolation error in this batch learning process. To mitigate this, Fujimoto, Meger, and Precup (2018) proposed batch-constrained deep Q-learning (BCQ) which uses

dual Q-networks, similar to TD3 (Fujimoto, van Hoof, & Meger, 2018), along with a generative model for action proposal and a perturbation model for constrained action adjustment. Algorithms specifically designed for learning off a fixed dataset, such as BCQ, would be required to practically and safely pretrain an RL controller from historic operating data.

5.4 Summary

The behaviour of a RL controller based on Lau's (2016) implementation the model-free off-policy DDPG algorithm (Lillicrap et al., 2015) was explored. This occurred through its interactions with a simulated open circuit hydrocyclone system based on Moly-Cop Tools CycloSim (Sepúlveda, 2012), which was extended to include roping criteria and a fixed size feed tank. The experimental design included industrially relevant considerations such as typically measured variable, restricted MV changes, safety limits, and the partial observability of the system through unmeasured variables and sensor noise. The controller was able to interact with the simulated system through three continuous MV; hydrocyclone feed solids concentration, hydrocyclone feed flow rate, and apex diameter. A reward function was constructed to guide the controller training process to bring the tank level and hydrocyclone underflow P80 to SP whilst avoid roping conditions and safety limits.

The simulated environment was made progressively more challenging through the inclusion additional sources of partial observability; unmeasured tank solids feed rate, unmeasured feed sizing %passing modifications, and simulated sensor noise. The inclusion of historic observations in the state representation was found to compensate well for the partial observability of the system. Though for the methods trialled the controller began to struggle under some of the more extreme sensor noise conditions. This being a good illustration that there should still be control override protections considered to compensate for atypical or adverse conditions.

The desirable process control behaviour of avoiding entering the roping fault state was learnt and apparent in testing. This would occur even when it resulted in the controller unable to reach SPs and instead operate in a region near the achievable

limit. The trained controller was also shown to exhibit the ability to generalise through being tested on initialisation conditions outside those experienced in training; tank level, feed sizing %passing modification, and unbalance tank inflow/outflow. The requirement for SP changes to be incorporated during the training process if they can or will be change in operation was demonstrated by degraded performance of the controller trained on fixed SP, even though SP difference formed part of the state representation. Considerations for the structure of the state representation to remain valid for SP changes that would occur in continuous process control were highlighted, and successful control following a number of SP changes was demonstrated.

A final discussion was also presented highlighting additional points of interest and potential limitations for RL based mineral processing circuit control. The frequency of state assessment and control changes, along with the extent to which MVs can adjusted at each time step, should be based on the intended control behaviour and an assessment of the system dynamics. Controller design would need to be aligned with considerations for more complex systems such as the incorporation of RNNs for longer-term temporal information, requirements to perform action selection or discrete actions, and appropriate filtering of sensor data. The requirement for development of methods to successfully pretrain RL controllers on historic operating data was also raised as a critical hurdle to be overcome for their safe and effective implementation in operating mineral processing circuits.

6 Conclusion

The core objectives of this project were to investigate the applicability of CNNs for hydrocyclone operating state detection and underflow particle size inference, through image-based monitoring of the underflow discharge. This was undertaken with a focus on industrial applicability, factors affecting performance, and ease of model training and deployment. Exploring the use RL for hydrocyclone control, incorporating outputs from the CNN sensors or a related competing method, was also the final core objective.

In addressing the hydrocyclone state detection objective, a proof of concept for the use of CNNs in monitoring images of hydrocyclone underflow discharge was demonstrated. Transfer learning using pretrained models to extract features allowed fixed camera three-state (fan/roping/blocked) classifiers to be trained on limited data, in both laboratory and industrial settings. The sensor development process and testing focused on industrially relevant considerations including image noise, lighting, transition operating states, image cropping, camera movement and obscuration.

The CNN based fixed camera system developed exhibits improvements in robustness to industrial conditions and simpler model development pipeline than previous image-based methods. This optical method offers a viable alternative to the commercially available hydrocyclone state detection options (such as acoustic, probe, or laser-based) in circuits where the underflow discharge is visible and unobstructed. This method also holds potential as a low-cost option for inhouse sensor development using a standard frame rate camera, sufficient lighting, and open-source software with included pretrained models. Further appealing attributes include no structural modifications required for sensor installation allowing for ease of retrofitting, and being a non-contact method the maintenance requirement would be minimised.

The training of a general hydrocyclone state detector that would not require site specific retraining may be possible, as indicated by the industrial non-fixed camera performance, but would require further investigation to confirm. Such a system would allow for rapid roll-out of fixed camera-based state detectors to new mineral processing sites and open the possibility of moving drone mounted sensors as proposed in Giglia and Aldrich (2020). It is likely that the training of such a model would require a significantly large dataset with sufficient variability in source examples, rather than taking many examples from a few sources, to help produce a model that can truly generalise.

In addressing the hydrocyclone underflow particle size estimation objective, a preliminary study into the use of CNN models to infer particle size information from video frames of hydrocyclone underflow discharge was undertaken. This being at an

individual hydrocyclone level for which online particle size monitoring options available to industry are limited, particularly for the underflow stream. A crop position near the hydrocyclone apex, and capturing the discharge stream edges, produced higher source footage classification accuracies than a crop position lower in the stream which captures a more disperse flow appearance. The inclusion of temporal information in the form of stacked consecutive frames did not produce a significant improvement in footage classification accuracy under the conditions tested and was not pursued further. Following this exploration of design aspects under which a fine-tuned CNN model could best allocate laboratory hydrocyclone high-speed video footage correctly, underflow particle size P80 modelling was performed and showed sufficient fine versus coarse discernment to warrant further investigation.

Under industrial conditions, and using a standard frame rate camera, image-based CNN fine-tuning and sensor ANN modelling methods were investigated. ANN based models relating underflow particle size P80 to hydrocyclone sensor information and sensor information of the wider grinding circuit were found to suffer in testing when sensor readings were outside that experienced during training. Both the image-based fine-tuned CNN model and the combined CNN model extension incorporating grinding circuit sensor information were able to outperform the sensor only ANN models. The combined model however did not show a significant testing performance improvement over the solely image-based model under these conditions. The long-term performance of a CNN based model in terms of handling atypical operating conditions, particle size extremes, and future variability is not clear from the current investigation, particularly given the narrow period of industrial data available.

In addressing hydrocyclone control using RL, high-level control behaviour on a simulated open circuit hydrocyclone was presented. A simple reward function considering PV to SP difference and additional penalties for undesirable conditions was sufficient to control this system under a range of conditions. This included avoiding conditions that were learnt to result in the roping hydrocyclone fault state and instead operating near that limit, even if the SPs thus could not be reached. SP

changes could be handled by the control only if varied SPs had formed part of the training process, making it an important design consideration if this behaviour is required.

Limits in allowable action step size and range led to bounded system changes as would be desirable for industrial control. Sensible control overrides, such as water addition for low tank level, were indicated as being useful to ensure safe and stable operation under industrial conditions. The inclusion of recent historic sensor information in the state representation was able to compensate for partial observability, such as unmeasured variables and sensor noise. As uncertainty increased, in this case through higher magnitude random noise, a longer history was required in the state representation to achieve satisfactory control performance.

This work supports the notion that as RL algorithms continue to develop they hold potential as a data-driven control method for more complex industrial systems, removing a level of subjectivity in the devising of process control behaviour. The ability to successfully pretrain the controller on historic data will be critical for its safe implementation in operating mineral processing plants. While the production of a complete process model may not be attainable, a combination of future novel algorithm designs and extensive detailed historical operating data would help to facilitate pretraining of model-free RL algorithms. The continued incorporation of additional sensors into modern mineral processing plant will also aid in defining the current state of the system, with novel sensors developed from techniques such as CNN based image sensors opening up more monitoring possibilities.

6.1 Recommendations

Considering the results yielded from the investigations undertaken, recommendations for future work on each objective are proposed. For objective one's image-based hydrocyclone operational state detection using CNNs, the following areas would be of interest

- The industrial fixed camera state detection model development should be repeated with an applied light source directed at the hydrocyclone underflow. This would be both to see if the daylight classification accuracy

can be improved with direct illumination and to assess the performance at night. If the presence of overexposed regions in the footage are found to significantly affect performance, then cost effective options to block direct sunlight from troublesome regions could be investigated.

- The long-term performance and stability of an industrial fixed camera system (both the state classifier and sensor health check) should be confirmed to be robust against all visually challenging conditions that may be encountered, including extreme weather events. The sensor health check tracker drift over an extended period should also be assessed. If drift is found to be an issue, then periodic correction should be undertaken such as recentring the tracker's filter with the initialisation frame as proposed by Bolme et al. (2010). It should be investigated if a drift compensation process can be successfully automated to remove the burden of regular manual checks and correction.
- Investigating the use of a thermal imaging camera in place of the visible light cameras already trialled would be of interest. Thermal camera's detection of heat energy would remove the requirement for an applied light source and provide the ability to perform detection through sources of obscuration such as steam (FLIR, n.d.) which could be beneficial in improving sensor performance. Considering the distinct slurry discharge characteristics of the three operating states, a lower resolution thermal image may be sufficient for classification. If successful, and found to be more robust, as prices of the thermal cameras decrease this may become the more desirable sensing option.
- Undertake further confirmation investigations into if water or dilute slurry draining from an industrial hydrocyclone could be used as a substitute for roping examples for model fine-tuning or testing of an existing model, as suggested in Giglia and Aldrich (2020).
- Attempt to build a large and varied dataset of industrial hydrocyclone discharges images: including a wide range of commodities, hydrocyclone systems, viewing angles, etc. This could provide an opportunity to fine-tune the convolutional layers, as suggested in Giglia and Aldrich (2020), with the

ultimate goal being to produce a model with strong enough generalisation performance to not require site specific fine-tuning for new installations.

- More generally the success of CNNs to detect the visible differences in hydrocyclone operating states opens the possibility of applying these techniques to assess the operation of other pieces of mineral processing equipment.

For objective two's image-based hydrocyclone underflow particle size estimation using CNNs, the following areas would be of interest

- Build a larger industrial dataset over an extend period, with the intent being to ensure a representative range of operating conditions is encountered. This would help to avoid extreme conditions and sensor readings which can deteriorate model performance.
- As only preliminary investigations were undertaken, further exploration of network architectures and extent of fine-tuning (facilitated by a larger dataset) would be of interest.
 - This could include considering temporal information using RNNs to incorporate past image features and/or sensor information. While the demand of these networks may affect frequency at which predictions are produced, if sufficiently beneficial for accuracy and considering the practical rate at which particle size information would be required, this may be deemed acceptable.
 - This could also include an investigation into the effect of the length of footage prior to sampling used in model training, and would be of most benefit performed under industrial conditions with a standard frame rate camera. This would be to see if an optimum can be found balancing the need to capture normal variability without being too long as to capture instability or changes in operation.
- If further success is found in producing an image-based sensor it is still worth assessing whether alternatives are suitable for a given site's requirements. If only circuit or cluster level particle size prediction are required, then it is worth investigating whether a soft sensor from available circuit sensor data

is sufficient. While attempts in this project were not successful, examples such as Zhang and Liang (2016) show there is potential to this method. Given there is no requirement to modify the hydrocyclone structure and the existing sensor data is readily available, this presents a minimal cost option to explore through collection of a large dataset. The use of vibration sensors in hydrocyclone particle size estimation systems, such as that indicated as being available from Emerson Electric (Cahill, 2021), would also be worth considering. This could present a lower cost of entry than an image-based sensor given the camera, lighting, and any other illumination control structures (to block out external lighting) that may be required.

For objective three's RL for hydrocyclone control the following areas would be of interest

- Further exploration of appropriate RL algorithms, neural network structures, reward function construction, controller interaction frequency and action size would all be beneficial in establishing a strong performing RL control. This should not just be considered in terms the simple open circuit hydrocyclone system explored in this work. Requirement for discrete actions for hydrocyclone on/off, material recirculation for closed circuit hydrocyclone systems, and other more complex mineral processing circuits should also be explored.
- The ability to handle sensor noise, uncertainty, and insufficient information leading to the partial observability of the system are key considerations for industrial control. Thus, the incorporation of RNNs and their use of temporal information to potentially counteract partially observability and other complexities in mineral processing systems would be of interest.
- A critical hurdle for applying RL control to mineral processing plants will be the ability to successfully pretrain on historic operating data, as the ability to freely explore the system will not be possible. Further investigation into algorithm design with this requirement in mind will be required to progress from simulation to applied operating circuit control.

- If simulated and lab/pilot scale RL control can be successfully pretrained and lead to sufficient control performance, then trialling control of a minor stream of an operating industrial plant would be a sensible next step. An example of this could be a regrind circuit, in which control could be trialled for the hydrocyclone cluster, potentially then expanded to include the mill performance, and then any subsequent processing units such as flotation cells. In this way troubleshooting and refining of the control can be performed on a stream with lower impact on production, whilst building confidence in the method on site.

Appendices

Appendix A

Multiple Source Industrial Hydrocyclone Image Datasets Frames Breakdown

Training Fan State Data Set		
Source	Footage ID	Frames
online	hydrocyclone_used_for_minerals_ores_classifying	87
donation	industry_C_cyc_A_1	35
donation	industry_C_cyc_B_1	35
donation	industry_C_cyc_B_4	35
donation	industry_C_cyc_C	35
donation	industry_B_1	19
donation	industry_B_2	19
donation	industry_B_4	19
donation	industry_B_5	19
donation	industry_B_6	19
<i>TOTAL</i>		322

Training Rope State Data Set		
Source	Footage ID	Frames
online	hydrocyclone_roping_incorrect_cyclone_operation	64
donation	industry_C_cyc_B_3	64
donation	industry_A_5	64
lab	underflow_2_roping	130
<i>TOTAL</i>		322

Validation Fan State Data Set		
Source	Footage ID	Frames
online	cyclone_output_from_evowash	35
online	hydrocyclone_used_for_minerals_ores_classifying	22
donation	industry_C_cyc_A_2	35
donation	industry_B_7	10
<i>TOTAL</i>		102

Validation Rope State Data Set		
Source	Footage ID	Frames
online	hydrocyclone_on_work	64
lab	underflow_2_roping	38
<i>TOTAL</i>		102

Note. From “Operational State Detection in Hydrocyclones with Convolutional Neural Networks and Transfer Learning,” by K. C. Giglia and C. Aldrich, 2020, *Minerals*

Engineering, 149, p. 8 (<https://doi.org/10.1016/j.mineng.2020.106211>). Copyright 2020 by Elsevier. Reprinted with permission.

Appendix B

Single Source Industrial Cyclone A Image Datasets Frames Breakdown

	Blocked	Fan	Rope
Train A	100	100	100
Train B	100	100	100
Val A	60	60	60
Test A	283	115	16
Test B	225	100	24
Test C	368	236	66
Test D	69	144	44
Test E	339	218	25

Appendix C

Three-State Confusion Matrices for Each Cyclone A Test Dataset and Crop Method

Test A Crop: Fixed		Predicted State		
		blocked	fan	rope
True State	blocked	178	10	95
	fan	1	114	0
	rope	4	1	11

Test A Crop: Random		Predicted State		
		blocked	fan	rope
True State	blocked	169	5	109
	fan	1	114	0
	rope	4	2	10

Test A Crop: Expanded		Predicted State		
		blocked	fan	rope
True State	blocked	153	3	127
	fan	31	83	1
	rope	4	1	11

Test B Crop: Fixed		Predicted State		
		blocked	fan	rope
True State	blocked	148	0	77
	fan	0	100	0
	rope	5	0	19

Test B Crop: Random		Predicted State		
		blocked	fan	rope
True State	blocked	147	0	78
	fan	3	97	0
	rope	3	0	21

Test B Crop: Expanded		Predicted State		
		blocked	fan	rope
True State	blocked	127	0	98
	fan	19	75	6
	rope	8	0	16

Test C Crop: Fixed		Predicted State		
		blocked	fan	rope
True State	blocked	255	10	103
	fan	0	236	0
	rope	4	0	62

Test C Crop: Random		Predicted State		
		blocked	fan	rope
True State	blocked	227	17	124
	fan	2	234	0
	rope	3	0	63

Test C Crop: Expanded		Predicted State		
		blocked	fan	rope
True State	blocked	168	52	148
	fan	15	214	7
	rope	4	2	60

Test D		Predicted State		
Crop: Fixed		blocked	fan	rope
True State	blocked	66	0	3
	fan	18	126	0
	rope	9	0	35

Test D		Predicted State		
Crop: Random		blocked	fan	rope
True State	blocked	62	0	7
	fan	23	121	0
	rope	6	0	38

Test D		Predicted State		
Crop: Expanded		blocked	fan	rope
True State	blocked	49	0	20
	fan	56	86	2
	rope	10	0	34

Test E		Predicted State		
Crop: Fixed		blocked	fan	rope
True State	blocked	221	31	87
	fan	0	218	0
	rope	0	0	25

Test C		Predicted State		
Crop: Random		blocked	fan	rope
True State	blocked	224	23	92
	fan	0	218	0
	rope	3	0	22

Test C		Predicted State		
Crop: Expanded		blocked	fan	rope
True State	blocked	222	10	107
	fan	1	210	7
	rope	3	0	22

Appendix D

Operating and Particle Size Information of Laboratory Experimental Runs

Sample ID	Feed Material	Feed %solids	Feed Pressure (kPa)	Underflow P80 (µm)	Underflow P50 (µm)
Q1_SL_P1	Quartz P80 = 748µm P50 = 371µm	3.3%	85	766	497
Q1_SL_P2		3.3%	75	675	428
Q1_SM_P1		5.5%	80	713	448
Q1_SM_P2		5.5%	70	646	399
Q1_SH_P1		7.2%	75	686	419
Q1_SH_P2		7.2%	60	661	403
Q2_SL_P1		Quartz P80 = 492µm P50 = 225µm	4.4%	80	551
Q2_SL_P2	4.4%		60	522	345
Q2_SM_P1	7.2%		80	515	327
Q2_SM_P2	5.3%		80	547	369
Q2_SM_P3	5.3%		60	527	349
Q2_SH_P1*	7.4%		80	522	345
Q2_SH_P2	7.4%		55	513	335
Q2_SH_P3*	7.4%		80	486	307

Note. Feed %solids are indicative only, based on initial water mass and solids addition. Thus, initial water held in pumping system, progressive removal of underflow samples, and slurry splashing were neglected in the calculation. *Samples were under similar operating conditions. During the initial sample (Q2_SH_P1) the underflow was unstable. Following the low pressure test the feed pressure was restored to 80kPa and as the underflow appeared more stable a second sample was taken (Q2_SH_P3).

Appendix E

Laboratory Footage Frames Used in Training and Validation Dataset Construction

Image Type	Sample ID	Underflow P80 (μm)	Training Frames	Validation Frames
rgb	Q2_SH_P3	486	1 - 8000	8001 - 10000
	Q2_SL_P2	522	1 - 8000	8001 - 10000
	Q2_SM_P3	527	1 - 8000	8001 - 10000
	Q2_SL_P1	551	2001 - 10000	11501 - 13500
	Q1_SM_P2	646	9001 - 17000	17001 - 19000
	Q1_SH_P2	661	1001 - 9000	9001 - 11000
	Q1_SL_P2	675	5001 - 13000	13001 - 15000
	Q1_SL_P1	766	7001 - 15000	15001 - 17000
rgb_motion	Q2_SH_P3	486	3 - 8002	8003 - 10002
	Q2_SL_P2	522	3 - 8002	8003 - 10002
	Q2_SM_P3	527	3 - 8002	8003 - 10002
	Q2_SL_P1	551	2001 - 10000	11501 - 13500
	Q1_SM_P2	646	9001 - 17000	17001 - 19000
	Q1_SH_P2	661	1001 - 9000	9001 - 11000
	Q1_SL_P2	675	5001 - 13000	13001 - 15000
	Q1_SL_P1	766	7001 - 15000	15001 - 17000

Appendix F

Laboratory Footage Frames Used in Test Dataset Construction

Image Type	Sample ID	Underflow P80 (μm)	Test Frames
rgb	Q2_SH_P2	513	13501 - 15500
	Q2_SM_P1	515	9601 - 11600
	Q2_SH_P1	522	11501 - 13500
	Q2_SM_P2	547	601 - 2600
	Q1_SH_P1	686	15601 - 17600
	Q1_SM_P1	713	5251 - 7250

Appendix G

CycloSim Data_File Spreadsheet Containing Supplied Base Case Data

Moly-Cop Tools™ (Version 3.0)

CYCLOSIM HYDROCYCLONE SIMULATOR

Simulation N°

Remarks

# of Cyclones	Cyclones Geometry, inches				
	Diameter	Height	Inlet	Vortex	Apex
<input type="text" value="10"/>	20.00	75.00	3.50	7.50	3.67
	<i>Default Values:</i>				
	60.00	5.00	7.00	3.75	

Operating Conditions : Ore Density, ton/m3

Cyclone Feed Conditions						
Dry Tons ton/hr	Water m3/hr	Slurry ton/hr	Flow m3/hr	Density ton/m3	% Solids (by volume)	% Solids (by weight)
1622.84	986.16	2609.01	1565.75	1.666	37.02	62.20

Classifier Constants : (from Cyclobal_Single, Ballbal_Direct or Ballbal_Reverse)

a1	a2	a3	a4	λ	Bp (coarse)	
9.680	1.401	54.964	0.523	0.950	0.050	
9.680	1.401	54.964	0.523	0.950	0.000	<i>Default Values</i>

Feed Size Distribution :

i	Mesh	Opening	Mid-Size	ton/hr	% Ret	% Pass
1	1.05	25400	30206	0.00	0.00	100.00
2	0.742	19050	21997	0.00	0.00	100.00
3	0.525	12700	15554	18.17	1.12	98.88
4	0.371	9500	10984	41.49	2.56	96.32
5	3	6700	7978	40.76	2.51	93.81
6	4	4750	5641	37.01	2.28	91.53
7	6	3350	3989	38.53	2.37	89.16
8	8	2360	2812	44.14	2.72	86.44
9	10	1700	2003	51.21	3.16	83.28
10	14	1180	1416	66.98	4.13	79.16
11	20	850	1001	78.83	4.86	74.30
12	28	600	714	100.61	6.20	68.10
13	35	425	505	122.86	7.57	60.53
14	48	300	357	142.11	8.76	51.77
15	65	212	252	143.62	8.85	42.92
16	100	150	178	125.60	7.74	35.18
17	150	106	126	98.97	6.10	29.08
18	200	75	89	73.25	4.51	24.57
19	270	53	63	53.99	3.33	21.24
20	400	38	45	38.98	2.40	18.84
21	-400	0	19	305.76	18.84	0.00
Total				<input type="text" value="1622.84"/>	<input type="text" value="100.00"/>	

Note. From "CycloSim Hydrocyclone Simulator in Moly-Cop Tools," by J. E. Sepúlveda, 2012, Molycop (<https://molycop.com/what-we-do/molycop-tools/>). Copyright 2012 by Moly-Cop. Reprinted with permission.

Appendix H

Copyright Permissions

Appendix H1 For Figure 2-1

6/7/20

Julius Kruttschnitt Mineral Research Centre
Isles Road, Indooroopilly, Queensland 4068, Australia

Dear Julius Kruttschnitt Mineral Research Centre,

It is my understanding that your organisation is the copyright holder for the following material:

Mineral comminution circuits: their operation and optimisation

T.J. Napier-Munn, S. Morrell, R.D. Morrison, and T. Kojovic

ISBN 0 646 28861 X

Figure 12.18 : The hydrocyclone, showing main components and principal flows (pg 310)

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

I would like to include the Figure 12.18 from Mineral comminution circuits: their operation and optimisation page 310 in the hydrocyclone related section of my thesis' literature review.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository space (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:**Mineral comminution circuits: their operation and optimisation****T.J. Napier-Munn, S. Morrell, R.D. Morrison, and T. Kojovic****ISBN 0 646 28861 X****Figure 12.18 : The hydrocyclone, showing main components and principal flows (pg 310)**

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository, subject to acknowledging the above reference in his thesis. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: 

Name: T.J. NAPIER-MUNN

Position: Emeritus Professor

Date: 14 July 2020

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H2 For Figure 2-3 and Figure 2-4

11/04/21

<https://github.com/ashushekar>
Dublin, Ireland

Dear Ashwin Naidu,

It is my understanding that you are the copyright holder for the following material:

**<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure3.jpg>
<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure4.jpg>**

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

I would like to include the two figures stated above in the convolutional neural networks related section of my thesis' literature review.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository espace (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:

<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure3.jpg>

<https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure4.jpg>

I hereby give permission for Keith Giglia to include the abovementioned materials in his higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.



Name: Ashwin Chandrashekar Naidu

Position: Data Scientist at Citi

Date: 12-April-2021

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H3 For Figure 2-5, Figure 2-6, and Figure 2-7

Reprinted under MIT license

<https://github.com/cs231n/cs231n.github.io/blob/master/LICENSE>

The MIT License (MIT)

Copyright (c) 2015 Andrej ~~Karpathy~~

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE

Appendix H4 For Figure 2-9

Adapted from Creative Commons CC0 1.0 Universal Public Domain Dedication file

https://commons.wikimedia.org/wiki/File:Reinforcement_learning_diagram.svg

I, the copyright holder of this work, hereby publish it under the following license:

This file is made available under the Creative Commons CC0 1.0 Universal Public Domain Dedication.

The person who associated a work with this deed has dedicated the work to the public domain by waiving all of their rights to the work worldwide under copyright law, including all related and neighboring rights, to the extent allowed by law. You can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission.

Appendix H5 Industry donated hydrocyclone underflow footage used in
Section 3.2

Appendix H5.1 BHP Nickel West, Mt Keith

17/04/18

125 St Georges Terrace,
Perth WA 6000

Dear Brent Sullivan (BHP, Mt Keith mine site)

It is my understanding that your organisation is the copyright holder for the following material:

Hydrocyclone video footage from BHP Nickel West Mt Keith mine site

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

The frames of the video footage in their entirety, which will be used for algorithm training but the images themselves will not be visible in the thesis or any subsequent publications.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository espace (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.]

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:

Hydrocyclone video footage from BHP Nickel West Mt Keith mine site (as per cover letter)

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his/her higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: 

Name: Brent Sullivan

Position: Senior Metallurgist

Date: 24/09/2018

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H5.2 IGO, Nova

31/01/20

Suite 4, Level 5, South Shore Centre,
85 South Perth Esplanade, South Perth WA 6151

Dear Matthew Spagnolo (IGO, Nova mine site)

It is my understanding that your organisation is the copyright holder for the following material:

Hydrocyclone video footage from Nova mine site

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

The frames of the video footage in their entirety, which may also be included in the PhD thesis body for illustrative purposes. The frames will not be made visible in any subsequent publications.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository espace (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:**Hydrocyclone video footage from Nova mine site (as per cover letter)**

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his/her higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: Name: Matthew SpagnoloPosition: Manager Processing & MaintenanceDate: 11/2/2020

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H5.3 Minara Resources, Murrin Murrin

19/06/18

Level 3, 30 The Esplanade,
Perth WA 6000

Dear Peter Hancock (Minara Resources, Murrin Murrin mine site)

It is my understanding that your organisation is the copyright holder for the following material:

Hydrocyclone video footage from Murrin Murrin mine site

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

The frames of the video footage in their entirety, which will be used for algorithm training but the images themselves will not be visible in the thesis or any subsequent publications.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository [espace \(http://espace.curtin.edu.au\)](http://espace.curtin.edu.au), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:

Hydrocyclone video footage from Murrin Murrin mine site (as per cover letter)

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his/her higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: 

Name:

Position:

Date:

Peter Hancock
Vice President - Minera Resources
21-06-2018

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H6 Anonymous industry donated hydrocyclone underflow footage used in Section 3.2



4 October 2019



Dear [REDACTED]

It is my understanding that your organisation is the copyright holder for the following material:

Hydrocyclone video footage from [REDACTED] mine site

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material/extract that I would like to use for the purposes of the thesis is:

The frames of the video footage in their entirety, which will be used for algorithm training but the images themselves will not be visible in the thesis or any subsequent publications.

Any reference to the data and its source will be anonymized, with a redacted version of this permission letter made available as evidence that the copyright owner's usage permission has been sought and granted.

Relevant sections of the thesis/publications employing this data will be made available for viewing by designated [REDACTED] staff to confirm anonymization is sufficient prior to publication.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository space (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely,

Keith Giglia (PhD student) and

Professor Chris Aldrich (Deputy Head: Western Australian School of Mines – Minerals, Energy and Chemical Engineering)

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:**Hydrocyclone video footage from [REDACTED] mine site (as per cover letter)**

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: [REDACTED]

Name: [REDACTED]

Position: [REDACTED]

Date: [REDACTED]

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H7 For Figure 4-7, Figure 4-8, Figure 4-13, Figure 4-40, and material from IGO's Nova mine site (videos, images, data, and information) referred to throughout Section 4.2 and captured by the dataset IGO Nova (2020).

IGO Nova. (2020). *Nova mine site data* [Data set, videos, and images].

Location: IGO Nova.

24/11/20

Suite 4, Level 5, South Shore Centre,
85 South Perth Esplanade, South Perth WA 6151

Dear Matthew Spagnolo (IGO, Nova mine site)

It is my understanding that your organisation is the copyright holder for the following material:

Hydrocyclone underflow video footage, particles sizing, circuit sensor information, and experimental setup images from Nova mine site.

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

The frames of the video footage in their entirety for algorithm training. The frames of the video footage and images of the experimental setup may also be included in the PhD thesis body and any additional publications (following approval by IGO staff).

Operational related information, including particle sizing, sensor data, and blend information, may be utilised for algorithm training and included in the PhD thesis body. The inclusion of any operational related information in any additional publications may only occur following approval by IGO staff.

IGO Nova mine site will be cited as the material source within the thesis, but the material source will remain anonymous in any additional publications. Relevant sections of the thesis/publications employing this material will be made available for viewing by designated IGO staff to confirm copyright permissions have been correctly implemented prior to publication.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository espace (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:

Hydrocyclone underflow video footage, particles sizing, circuit sensor information, and experimental setup images from Nova mine site (as per cover letter)

I hereby give permission for Keith Giglia to include the abovementioned materials in his higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed:



Name: Matthew Spagnolo

Position: Manager Processing, IGO Nova

Date: 12/12/20

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H8 For Appendix G and the Moly-Cop Tools CycloSim (CycloSim_Single.xlsx) software and supplied base case example data used as the basis for the simulated hydrocyclone circuit in Section 5

16/04/18

Ave Pedro de Valdivia, 0168
Santiago, Chile

Dear Darren O'Connell (Moly-Cop Chile S. A)

It is my understanding that your organisation is the copyright holder for the following material:

Moly-Cop Tools™ 3.0

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. The subject of my research is applying machine learning techniques to the monitoring and control of hydrocyclones. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

The CycloSim_Single simulation calculations and supplied base case data to produce a simulated environment in the Python programming language for the machine learning control to interact with. My thesis and any publications will refer to the fact that this simulated environment was based on CycloSim_Single and the supplied base case. I would also include the pages (text only, not active Excel formulas) of CycloSim_Single with supplied base case data as an appendix in my thesis for transparency.

Once completed, the thesis will be made available in online form via Curtin University's Institutional Repository espace (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

If you are not the copyright owner of the material in question, I would be grateful for any information you can provide as to who is likely to hold the copyright.

I look forward to hearing from you and thank you in advance for your consideration of my request.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:**Moly-Cop Tools™ 3.0 CycloSim_Single** (as per cover letter)

I hereby give permission for Keith Giglia to include the abovementioned material(s) in his/her higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed:



Name: Darren O'Connell

Position: VP Marketing

Date: 24/4/18

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H9 For Ben Lau's DDPG-Keras-Torcs software which was modified for use in Section 5

From: Ben Lau
Sent: Tuesday, 26 March 2019 7:12 PM
To: Keith Giglia
Subject: Re: DDPG-Keras-Torcs Permission

Hi, The code is meant to be open to public. Feel Free to use for your own research purpose. I just hope that you can cite my work.

Thanks.

Ben

Keith Giglia <keith.giglia@postgrad.curtin.edu.au> 於 2019 年 3 月 26 日週二 下午 12:40 寫道 :

Hi Ben,

My name is Keith Giglia and am a PhD student at Curtin University, Western Australia, currently researching into applications of machine learning into process monitoring and control.

Firstly, I just wanted to thank you for your post & code related to DDPG with Keras for Torc, it was great for someone coming from a non-computer science background to wrap my head around the theory and use.

I'm contacting you as I would like to use your code related to DDPG-Keras-Torcs as the basis for my own implementation as part of my research. If this is acceptable could you please fill in the attached copyright permissions letter, and know that both my thesis and any resulting publications will reference your implementation as forming the basis of that used in the research.

Thank you for considering my request.

Regards

Keith

Appendix H10 For material from co-authored publication; predominantly in Section 2.3.4, Section 3.1, Section 3.2, and Appendix A. Includes Figure 2-2, Figure 3-1, Figure 3-2, Figure 3-4, Figure 3-5, Figure 3-6, Figure 3-8, Table 3-1, Table 3-2, and Table 3-3

<https://doi.org/10.1016/j.mineng.2020.106211>

This article was published in *Minerals Engineering*, Vol 149, K. C. Giglia & C. Aldrich, Operational state detection in hydrocyclones with convolutional neural networks and transfer learning, Article 106211, pp. 1–9, Copyright Elsevier (2020).

Appendix H10.1 Co-author permission and attribution statement

26/09/22

Curtin University
Kent St, Bentley WA 6102

Dear Professor Chris Aldrich,

In relation to the co-authorship and copyright holding for the following material:

Giglia, K. C., & Aldrich, C. (2020). Operational state detection in hydrocyclones with convolutional neural networks and transfer learning. *Minerals Engineering*, 149, Article 106211. <https://doi.org/10.1016/j.mineng.2020.106211>

I would like to reproduce an extract of this work in a doctoral thesis which I am currently undertaking at Curtin University in Perth, Western Australia. I am carrying out this research in my own right and have no association with any commercial organisation or sponsor.

The specific material / extract that I would like to use for the purposes of the thesis is:

Text, figures and tables from the co-authored publication.

The attribution statement regarding the agreed respective contributions to the publication will also be included as below.

	Conception and Design	Acquisition of Data & Method	Data Conditioning & Manipulation	Analysis & Statistical Method	Interpretation & Discussion	Final Approval
Keith Giglia	X	X	X	X	X	X
I acknowledge that these represent my contribution to the above research output Signed: [Redacted]						
Prof Chris Aldrich	X					X
I acknowledge that these represent my contribution to the above research output Signed: [Redacted]						

Once completed, the thesis will be made available in online form via Curtin University’s Institutional Repository space (<http://espace.curtin.edu.au>), as per University requirements associated with my PhD program.

I would be most grateful for your consent to the copying and communication of the work as proposed. If you are willing to grant this consent, please complete and sign the attached approval slip and return it to me at the address shown. Full acknowledgement of the ownership of the copyright and the source of the material will be provided with the material.

Yours sincerely

Keith Giglia

PERMISSION TO USE COPYRIGHT MATERIAL AS SPECIFIED BELOW:

Text, figures and tables from the co-authored publication (as per cover letter)

I hereby give permission for Keith Giglia to include the abovementioned materials in his higher degree thesis for Curtin University, and to communicate this material via the espace institutional repository. This permission is granted on a non-exclusive basis and for an indefinite period.

I confirm that I am the copyright owner of the specified material.

Signed: 

Name: Chris Aldrich

Position: Professor

Date: 3 September 2022

Please return signed form to Keith Giglia (keith.giglia@student.curtin.edu.au)

Appendix H10.2 Publisher permission

Pages 227–229 have been redacted for confidentiality reasons.

Appendix H10.3 Publisher figure and table adaption permission

Pages 230–232 have been redacted for confidentiality reasons.

References

- 911 Metallurgy Corp. (2016, November 27). *Hydrocyclone roping incorrect cyclone operation* [Video]. YouTube.
<https://www.youtube.com/watch?v=nq908po87F0>
- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jozefowicz, R., Jia, Y., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Schuster, M., Monga, R., Moore, S., Murray, D., Olah, C., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., & Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems* (Version 1.8.0) [Computer software]. <https://www.tensorflow.org>
- Aldrich, C., Uahengo, F. D. L., & Kistner, M. (2015). Estimation of particle size in hydrocyclone underflow streams by use of multivariate image analysis. *Minerals Engineering*, *70*, 14-19.
<https://doi.org/10.1016/j.mineng.2014.08.018>
- Bascur, O. (2019). Process control and operational intelligence. In R. C. Dunne, S. K. Kawatra, & C. A. Young (Eds.), *SME mineral processing & extractive metallurgy handbook* (pp. 277-316). Society for Mining, Metallurgy, and Exploration.
- Bolme, D. S., Beveridge, J. R., Draper, B. A., & Lui, Y. M. (2010). Visual object tracking using adaptive correlation filters. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2544-2550.
<https://doi.org/10.1109/CVPR.2010.5539960>

- Botha, S., le Roux, J. D., & Craig, I. K. (2018). Hybrid non-linear model predictive control of a run-of-mine ore grinding mill circuit. *Minerals Engineering*, 123, 49-62. <https://doi.org/10.1016/j.mineng.2018.04.016>
- Bowers, S. V., Bassett, T. S., Banerjee, T., Schaffer, M., & Nower, D. (2019). *Monitoring and controlling hydrocyclones using vibration data* (U.S. Patent No. 10,258,905 B2). U.S. Patent and Trademark Office. <https://patentimages.storage.googleapis.com/a0/c8/0d/d4b69343cece3f/US10258905.pdf>
- Bradski, G. (2000). *The OpenCV library* Dr. Dobb's Journal of Software Tools [Computer software]. <http://www.drdobbs.com/open-source/the-opencv-library/184404319>
- Buttler, B. M., Bochicchio, M. P., Culbertson, J. R., & Singer, R. B. (2019). *Wireless hydrocyclone roping and wear management system* (U.S. Patent Application No. 2019/0176167 A1). U.S. Patent and Trademark Office. <https://patentimages.storage.googleapis.com/9c/fa/83/65c99f480b5fb2/US20190176167A1.pdf>
- Cahill, J. (2021, March 15). Optimizing hydrocyclone separation processes. *Emerson Automation Experts Blog*. <https://www.emersonautomationexperts.com/2021/industry/metals-mining-minerals/optimizing-hydrocyclone-separation-processes/>
- CDE Group. (2011, November 2). *Cyclone output from evowash* [Video]. YouTube. https://www.youtube.com/watch?v=8ozTlijUe_I

Chicco, D. (2017). Ten quick tips for machine learning in computational biology. *BioData Mining*, 10, Article 35. <https://doi.org/10.1186/s13040-017-0155-3>

Chollet, F. (2018). *Deep learning with python*. Manning Publications.

Chollet, F., & others. (2015a). *Keras* (Version 2.1.5-tf) [Computer software]. <https://keras.io>

Chollet, F., & others. (2015b). *Keras FAQ*. Keras. https://keras.io/getting_started/faq/

CiDRA. (2018). *CYCLONEtrac OSM: Oversize Monitoring System* [Brochure]. https://www.cidra.com/sites/default/files/document_library/BI0676-CYCLONEtrac-OSM-Gen2-Data-Sheet.pdf

CiDRA. (2019). *CYCLONEtrac PST: Particle size tracking system* [Brochure]. https://www.cidra.com/sites/default/files/document_library/BI0509-CYCLONEtrac-PST-Data-Sheet.pdf

Clark, A., & contributors. (2010). *Pillow (PIL fork)* (Version 5.3.0) [Computer software]. <https://python-pillow.org/>

Clark, A., & contributors. (n.d.). *Concepts*. Pillow. <https://pillow.readthedocs.io/en/stable/handbook/concepts.html#filters>

Coetzee, L. C. (2014). Stabilising and optimising a primary closed-loop milling circuit feeding a flotation circuit using StarCS RNMPC. *IFAC Proceedings Volumes*, 47(3), 9786-9791. <https://doi.org/10.3182/20140824-6-ZA-1003.01396>

- Coetzee, L. C., Craig, I. K., & Kerrigan, E. C. (2010). Robust nonlinear model predictive control of a run-of-mine ore milling circuit. *IEEE Transactions on Control Systems Technology*, 18(1), 222-229. <https://doi.org/10.1109/TCST.2009.2014641>
- Coetzee, L. C., & Ramonotsi, M. (2016). Applying StarCS RNMPC with real-time optimiser to Pilanesberg platinum mines primary UG2 milling circuit. *IFAC-PapersOnLine*, 49(20), 78-83. <https://doi.org/10.1016/j.ifacol.2016.10.100>
- Conradie, A. V. E., & Aldrich, C. (2001). Neurocontrol of a ball mill grinding circuit using evolutionary reinforcement learning. *Minerals Engineering*, 14(10), 1277-1294. [https://doi.org/10.1016/S0892-6875\(01\)00144-3](https://doi.org/10.1016/S0892-6875(01)00144-3)
- Cyclone Engineering Jinyang. (2016, April 21). *Hydrocyclone on work* [Video]. YouTube. <https://www.youtube.com/watch?v=Hy7pbzDUuc>
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248-255. <https://doi.org/10.1109/CVPR.2009.5206848>
- Draper, N. R., & Smith, H. (1998). *Applied regression analysis* (3rd ed.). John Wiley & Sons.
- Dubey, R. K., Climent, E., Banerjee, C., & Majumder, A. K. (2016). Performance monitoring of a hydrocyclone based on underflow discharge angle. *International Journal of Mineral Processing*, 154, 41-52. <https://doi.org/10.1016/j.minpro.2016.07.002>

Eren, H., Fung, C. C., & Gupta, A. (1996, March 24–28). *Application of artificial neural network in estimation of hydrocyclone parameters* [Paper presentation]. The AusIMM Annual Conference, Perth, WA, Australia.

Eren, H., Fung, C. C., & Wong, K. W. (1997). An application of artificial neural network for prediction of densities and particle size distributions in mineral processing industry. *IEEE Instrumentation and Measurement Technology Conference Sensing, Processing, Networking: IMTC Proceedings, 2*, 1118-1121.
<https://doi.org/10.1109/IMTC.1997.612374>

Eren, H., Fung, C. C., Wong, K. W., & Gupta, A. (1997). Artificial neural networks in estimation of hydrocyclone parameter d50c with unusual input variables. *IEEE Transactions on Instrumentation and Measurement*, 46(4), 908-912.
<https://doi.org/10.1109/19.650798>

Flintoff, B., & Knorr, B. (2019). Cyclones. In R. C. Dunne, S. K. Kawatra, & C. A. Young (Eds.), *SME mineral processing & extractive metallurgy handbook* (pp. 595-610). Society for Mining, Metallurgy, and Exploration.

Flintoff, B. C., Plitt, L. R., & Turak, A. A. (1987). Cyclone modelling: A review of present technology. *CIM Bulletin*, 80(905), 39-50.

FLIR. (n.d.). *Thermal imaging for automation / process control: Discover a wide variety of applications* [Brochure].
http://www.flirmedia.com/MMC/THG/Brochures/T820485/T820485_EN.pdf

FLSmidth. (2021). *SmartCyclone: Wireless sensor technology* [Brochure].
<https://flsmidth-prod-cdn.azureedge.net/-/media/brochures/brochures->

[products/centrifugation-and-classification/2020/krebs-smartcyclone-technical-datasheet.pdf?rev=f9a6bb2c-b6a7-4b81-b924-e33fbb4d7b97](https://doi.org/10.1016/j.mineng.2020.106211)

Fu, Y., & Aldrich, C. (2018). Froth image analysis by use of transfer learning and convolutional neural networks. *Minerals Engineering*, 115, 68-78.
<https://doi.org/10.1016/j.mineng.2017.10.005>

Fu, Y., & Aldrich, C. (2019). Flotation froth image recognition with convolutional neural networks. *Minerals Engineering*, 132, 183-190.
<https://doi.org/10.1016/j.mineng.2018.12.011>

Fujimoto, S., Meger, D., & Precup, D. (2018). *Off-policy deep reinforcement learning without exploration*. ArXiv. <https://doi.org/10.48550/arXiv.1812.02900>

Fujimoto, S., van Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. In J. Dy & A. Krause (Eds.), *Proceedings of the 35th international conference on machine learning* (Vol. 80, pp. 1587-1596). PMLR. <http://proceedings.mlr.press/v80/fujimoto18a/fujimoto18a.pdf>

Giglia, K. C., & Aldrich, C. (2020). Operational state detection in hydrocyclones with convolutional neural networks and transfer learning. *Minerals Engineering*, 149, Article 106211. <https://doi.org/10.1016/j.mineng.2020.106211>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., & Bengio, Y. (2013). *An empirical investigation of catastrophic forgetting in gradient-based neural networks*. ArXiv. <https://doi.org/10.48550/arXiv.1312.6211>

- Govindhasamy, J. J., McLoone, S. F., Irwin, G. W., French, J. J., & Doyle, R. P. (2005). Reinforcement learning for online industrial process control. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 9(1), 23-30. <https://doi.org/10.20965/jaciii.2005.p0023>
- Gu, S., Lillicrap, T., Sutskever, I., & Levine, S. (2016). *Continuous deep Q-learning with model-based acceleration*. ArXiv. <https://doi.org/10.48550/arXiv.1603.00748>
- Guo, L., Wang, H., & Zhang, J. (2019). Data-Driven grinding control using reinforcement learning. In Z. Xiao, L. T. Yang, P. Balaji, T. Li, K. Li, & A. Zomaya (Eds.), *IEEE 21st international conference on high performance computing and communications; IEEE 17th international conference on smart city; IEEE 5th international conference on data science and systems (HPCC/SmartCity/DSS)* (pp. 2817-2824). IEEE. <https://doi.org/10.1109/HPCC/SmartCity/DSS.2019.00395>
- Gupta, A., & Eren, H. (1990). Mathematical modelling and on-line control of hydrocyclones. *The AusIMM Proceedings*, 295(1), 31-41.
- Gupta, A., & Yan, D. S. (2006). *Mineral processing design and operation: An introduction*. Elsevier. <https://doi.org/10.1016/B978-0-444-51636-7.X5000-1>
- Gutiérrez, J. A., Dyakowski, T., Beck, M. S., & Williams, R. A. (2000). Using electrical impedance tomography for controlling hydrocyclone underflow discharge. *Powder Technology*, 108(2), 180-184. [https://doi.org/10.1016/S0032-5910\(99\)00218-1](https://doi.org/10.1016/S0032-5910(99)00218-1)

- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., & Levine, S. (2018). *Soft actor-critic algorithms and applications*. ArXiv. <https://doi.org/10.48550/arXiv.1812.05905>
- Hafner, R., & Riedmiller, M. (2011). Reinforcement learning in feedback control: Challenges and benchmarks from technical process control. *Machine Learning*, 84, 137-169. <https://doi.org/10.1007/s10994-011-5235-x>
- Hallén, M., Åstrand, M., Sikström, J., & Servin, M. (2019). Reinforcement learning for grinding circuit control in mineral processing. *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation*, 488-495. <https://doi.org/10.1109/ETFA.2019.8869212>
- Hausknecht, M., & Stone, P. (2015). *Deep reinforcement learning in parameterized action space*. ArXiv. <https://doi.org/10.48550/arXiv.1511.04143>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). *Deep residual learning for image recognition*. ArXiv. <https://doi.org/10.48550/arXiv.1512.03385>
- Heinisuo, O.-P. [skvark]., andrey.senyaev, asmorkalov, & sergregory. (2019). *opencv-python* (Version 4.1.2.30) [Computer software]. <https://pypi.org/project/opencv-python/4.1.2.30/>
- Heinisuo, O.-P. [skvark]., andrey.senyaev, asmorkalov, & sergregory. (2020). *opencv-contrib-python* (Version 4.2.0.32) [Computer software]. <https://pypi.org/project/opencv-contrib-python/4.2.0.32/>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>

- Hoskins, J. C., & Himmelblau, D. M. (1992). Process control via artificial neural networks and reinforcement learning. *Computers and Chemical Engineering*, 16(4), 241-251. [https://doi.org/10.1016/0098-1354\(92\)80045-B](https://doi.org/10.1016/0098-1354(92)80045-B)
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). *MobileNets: Efficient convolutional neural networks for mobile vision applications*. ArXiv. <https://doi.org/10.48550/arXiv.1704.04861>
- Hulbert, D. G. (1992). *Hydro-cyclone underflow monitor based on underflow slurry stream shape* (U.S. Patent No. 5,132,024). U.S. Patent and Trademark Office. <https://patentimages.storage.googleapis.com/e2/ab/00/4c9f6614b691cb/US5132024.pdf>
- Hulbert, D. G. (1993). *Method and apparatus for measuring shade of hydrocyclone underflow* (U.S. Patent No. 5,248,442). U.S. Patent and Trademark Office. <https://patentimages.storage.googleapis.com/41/f1/fd/cd2969d7ecc0a0/US5248442.pdf>
- IGO Nova. (2020). *Nova mine site data* [Data set, videos, and images]. Location: IGO Nova.
- Ioffe, S., & Szegedy, C. (2015). *Batch normalization: Accelerating deep network training by reducing internal covariate shift*. ArXiv. <https://doi.org/10.48550/arXiv.1502.03167>

- Janse van Vuuren, M. J., Aldrich, C., & Auret, L. (2011). Detecting changes in the operational states of hydrocyclones. *Minerals Engineering*, 24(14), 1532-1544. <https://doi.org/10.1016/j.mineng.2011.08.002>
- Janse van Vuuren, M. J., Aldrich, C., Auret, L., Bezuidenhoudt, C., & De Jager, C. (2010). On-line monitoring of hydrocyclones by use of image analysis. *IFAC Proceedings Volumes*, 43(9), 87-91. <https://doi.org/10.3182/20100802-3-ZA-2014.00021>
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., & Darrell, T. (2014). *Caffe: Convolutional architecture for fast feature embedding* [Computer software]. <https://arxiv.org/abs/1408.5093>
- Jiang, Y., Fan, J., Chai, T., Li, J., & Lewis, F. L. (2018). Data-driven flotation industrial process operational optimal control based on reinforcement learning. *IEEE Transactions on Industrial Informatics*, 14(5), 1974-1989. <https://doi.org/10.1109/TII.2017.2761852>
- Karpathy, A. [karpathy]. (2015a). *Maxpool*. GitHub. <https://github.com/cs231n/cs231n.github.io/blob/master/assets/cnn/maxpool.jpeg>
- Karpathy, A. [karpathy]. (2015b). *Neural_net2*. GitHub. https://github.com/cs231n/cs231n.github.io/blob/master/assets/nn1/neural_net2.jpeg

Karpathy, A. [karpathy]. (2015c). *Pool*. GitHub.

<https://github.com/cs231n/cs231n.github.io/blob/master/assets/cnn/pool.jpg>

Karpathy, A. [karpathy]., omoindrot, hammer, ranjaykrishna, alessandro-gentilini, undefdev, sameermanek, qxcv, nbro, michaelwu, moroshko, mbosnjak, kchen92, cjoshmartin, jin14, joadinissf, FCH808, crizCraig, beckgom, ccwang002, & Vincibean. (2015). *Convolutional neural networks (CNNs / ConvNets)*. GitHub.

<https://github.com/cs231n/cs231n.github.io/blob/master/convolutional-networks.md>

Kingma, D. P., & Ba, L. J. (2014). *Adam: A method for stochastic optimization*. ArXiv.

<https://doi.org/10.48550/arXiv.1412.6980>

Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 25, pp. 1097-1105). Curran Associates.

<https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>

Lau, B. [yanpanlau]. (2016). *DDPG-Keras-Torcs* [Computer software]. GitHub.

<https://github.com/yanpanlau/DDPG-Keras-Torcs>

le Roux, J. D., Olivier, L. E., Naidoo, M. A., Padhi, R., & Craig, I. K. (2016). Throughput and product quality control for a grinding mill circuit using non-linear MPC.

Journal of Process Control, 42, 35-50.

<https://doi.org/10.1016/j.jprocont.2016.04.007>

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436-444.

<https://doi.org/10.1038/nature14539>

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2015). *Continuous control with deep reinforcement learning*. ArXiv.

<https://doi.org/10.48550/arXiv.1509.02971>

Lin, B., & Jørgensen, S. B. (2011). Soft sensor design by multivariate fusion of image features and process measurements. *Journal of Process Control*, 21(4), 547-

553. <https://doi.org/10.1016/j.jprocont.2011.01.006>

Lin, M., Chen, Q., & Yan, S. (2013). *Network in network*. ArXiv.

<https://doi.org/10.48550/arXiv.1312.4400>

Lukežič, A., Vojíř, T., Čehovin Zajc, L., Matas, J., & Kristan, M. (2018). Discriminative correlation filter tracker with channel and spatial Reliability. *International Journal of Computer Vision*, 126(7), 671-688.

<https://doi.org/10.1007/s11263-017-1061-3>

Maron, R., O'Keefe, C., & Sepulveda, J. (2018). Assessing the benefits of automatic grinding control using PST technology for true on-line particle size measurement. *IOP Conference Series: Materials Science and Engineering*,

427, Article 012035. <https://doi.org/10.1088/1757-899x/427/1/012035>

- Matthews, B. W. (1975). Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2), 442-451. [https://doi.org/10.1016/0005-2795\(75\)90109-9](https://doi.org/10.1016/0005-2795(75)90109-9)
- Megajoice. (2017). *File:Reinforcement learning diagram.svg*. Wikimedia Commons. https://commons.wikimedia.org/wiki/File:Reinforcement_learning_diagram.svg
- Metso Outotec. (2020). *PSI 500i particle size analyzer*. <https://www.mogroup.com/portfolio/psi-500i-particle-size-analyzer/>
- Metso Outotec. (2021). *CycloneSense: Enabling optimization through sense* [Brochure]. <https://www.mogroup.com/globalassets/saleshub/documents--episerver/cyclonesense-enabling-optimization-through-sense-web.pdf>
- Mintek. (2011a). *CyLas: Keeping an eye on your cyclone* [Brochure]. <http://www.mintek.co.za/wp-content/uploads/2011/09/CyLas-brochure-en.pdf>
- Mintek. (2011b). *Milling Control & Optimisation: MillStar* [Brochure]. <http://www.mintek.co.za/wp-content/uploads/2011/09/MillStar-Brochure2.pdf>
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., & Kavukcuoglu, K. (2016). *Asynchronous methods for deep reinforcement learning*. ArXiv. <https://doi.org/10.48550/arXiv.1602.01783>
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C.,

- Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533. <https://doi.org/10.1038/nature14236>
- Moriarty, D. E., & Miikkulainen, R. (1996). Efficient reinforcement learning through symbiotic evolution. *Machine Learning*, 22, 11-32. <https://doi.org/10.1007/BF00114722>
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. MIT Press.
- Nageswararao, K. (1995). Technical note: A generalised model for hydrocyclone classifiers. *The AusIMM Proceedings*, 300(2), 21.
- Naidu, A. (2019a). *Figure3*. GitHub. <https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure3.jpg>
- Naidu, A. (2019b). *Figure4*. GitHub. <https://github.com/ashushekar/image-convolution-from-scratch/blob/master/images/figure4.jpg>
- Napier-Munn, T. J., Morrell, S., Morrison, R. D., & Kojovic, T. (2005). *Mineral comminution circuits: Their operation and optimisation*. Julius Kruttschnitt Mineral Research Centre.
- Napier, L. F. A., & Aldrich, C. (2017). An IsaMill soft sensor based on random forests and principal component analysis. *IFAC-PapersOnLine*, 50(1), 1175-1180. <https://doi.org/10.1016/j.ifacol.2017.08.270>
- Narasimha, M., Mainza, A. N., Holtham, P. N., Powell, M. S., & Brennan, M. S. (2014). A semi-mechanistic model of hydrocyclones — Developed from industrial

data and inputs from CFD. *International Journal of Mineral Processing*, 133, 1-12. <https://doi.org/10.1016/j.minpro.2014.08.006>

Neesse, T., Schneider, M., Golyk, V., & Tiefel, H. (2004). Measuring the operating state of the hydrocyclone. *Minerals Engineering*, 17(5), 697-703. <https://doi.org/10.1016/j.mineng.2004.01.015>

Ng, A. Y., Harada, D., & Russell, S. (1999, June 27–30). *Policy invariance under reward transformations: Theory and application to reward shaping* [Paper presentation]. Sixteenth International Conference on Machine Learning, Bled, Slovenia. <https://people.eecs.berkeley.edu/~russell/papers/icml99-shaping.pdf>

NVIDIA Corporation. (2014). *DIGITS* (Version 6.1.1) [Computer software]. <https://docs.nvidia.com/deeplearning/digits/>

O'Keefe, C., Rothman, P., Maron, R., Newton, D., Mercuri, J., Cirulis, D., & Holdsworth, M. (2014, April 7–10). *Grind circuit optimization at Rio Tinto Kennecott using real-time measurement of individual hydrocyclone overflow stream particle size enabled by novel CYCLONEtrac technology* [Paper presentation]. Comminution '14, Cape Town, South Africa. https://www.cidra.com/sites/default/files/document_library/BI0614-CiDRA-Comminution-2014-Ver20140321v3-Color.pdf

Oliphant, T. E. (2006). *Guide to NumPy* [Computer software]. <https://archive.org/details/NumPyBook>

- Olivier, J., & Aldrich, C. (2021). Underflow particle size estimation of hydrocyclones by use of transfer learning with convolutional neural networks. *IFAC-PapersOnLine*, 54(11), 85-90. <https://doi.org/10.1016/j.ifacol.2021.10.055>
- Olson, T. J., & Waterman, R. J. (2006). *Hydrocyclone roping detector and method* (U.S. Patent No. 6,983,850 B2). U.S. Patent and Trademark Office. <https://patentimages.storage.googleapis.com/3f/69/1e/a332e35df75315/US6983850.pdf>
- Orway Mineral Consultants. (2017, October 10). Cyclones don't do any grinding. *Orway Mineral Consultants Blog*. <https://orway.com.au/blog/cyclones-dont-grinding/>
- Outotec. (2009). *PSI 300: On-line slurry particle size analyzer* [Brochure].
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). *Scikit-learn: Machine learning in python* [Computer software]. *Journal of Machine Learning Research*, 12, 2825-2830. <https://scikit-learn.org/>
- Petersen, K. R. P., Aldrich, C., van Deventer, J. S. J., McInnes, C., & Stange, W. W. (1996). Hydrocyclone underflow monitoring using image processing methods. *Minerals Engineering*, 9(3), 301-315. [https://doi.org/10.1016/0892-6875\(96\)00015-5](https://doi.org/10.1016/0892-6875(96)00015-5)

- Platt, J. C. (2000). Probabilities for SV machines. In A. J. Smola, P. L. Bartlett, B. Schölkopf, & D. Schuurmans (Eds.), *Advances in large margin classifiers* (pp. 61-74). MIT Press.
- Plitt, L. R. (1976). A mathematical model of the hydrocyclone classifier. *CIM Bulletin*, 69(776), 114-123.
- Pueo, B. (2016). High speed cameras for motion analysis in sports science. *Journal of Human Sport and Exercise*, 11(1), 53-73.
<https://doi.org/10.14198/jhse.2016.111.05>
- Putz De La Fuente, E. (2019). *Hydrocyclone monitoring system and method* (International Publication No. 2019/173874 A1). World Intellectual Property Organization.
<https://patentimages.storage.googleapis.com/4e/ba/7c/b6ab19a6ba036c/WO2019173874A1.pdf>
- Razavian, A. S., Azizpour, H., Sullivan, J., & Carlsson, S. (2014). *CNN features off-the-shelf: An astounding baseline for recognition*. ArXiv.
<https://doi.org/10.48550/arXiv.1403.6382>
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211-252. <https://doi.org/10.1007/S11263-015-0816-Y>
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). *Proximal policy optimization algorithms*. ArXiv. <https://doi.org/10.48550/arXiv.1707.06347>

Sepúlveda, J. E. (2012). *Molycop tools* (Version 3.0) [Computer software]. Molycop.

<https://molycop.com/what-we-do/molycop-tools/>

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.

<https://doi.org/10.1038/nature16961>

Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., & Riedmiller, M. (2014). Deterministic policy gradient algorithms. In E. P. Xing & T. Jebara (Eds.), *Proceedings of the 31st International Conference on Machine Learning* (Vol. 32, pp. 387-395). PMLR. <http://proceedings.mlr.press/v32/silver14.pdf>

Simonyan, K., & Zisserman, A. (2014). *Very deep convolutional networks for large-scale image recognition*. ArXiv. <https://doi.org/10.48550/arXiv.1409.1556>

Spielberg, S. P. K. (2017). *Deep reinforcement learning approaches for process control* [Master's thesis, University of British Columbia]. UBC Theses and Dissertations.

<https://open.library.ubc.ca/collections/ubctheses/24/items/1.0361156>

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929-1958.

<https://dl.acm.org/doi/pdf/10.5555/2627435.2670313>

- Stange, W. (1992). The control of mineral processing plants using neural network techniques. *IFAC Proceedings Volumes*, 25(17), 13-18.
<https://doi.org/10.1016/B978-0-08-041704-2.50011-X>
- Strudwicke, C. D., Penna, W. D., & Baker, B. M. (2017). *Hydrocyclone system* (International Publication No. 2017/197451 A1). World Intellectual Property Organization.
<https://patentimages.storage.googleapis.com/d2/45/11/d560b11bd829d8/WO2017197451A1.pdf>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). The MIT Press.
- Syafiie, S., Tadeo, F., & Martinez, E. (2008). Model-free learning control of chemical processes. In C. Weber, M. Elshaw, & N. M. Mayer (Eds.), *Reinforcement learning* (pp. 295-310). IntechOpen. <https://doi.org/10.5772/5287>
- Szegedy, C., Wei, L., Yangqing, J., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1-9. <https://doi.org/10.1109/CVPR.2015.7298594>
- Tallec, C., Blier, L., & Ollivier, Y. (2019). *Making deep Q-learning methods robust to time discretization*. ArXiv. <https://doi.org/10.48550/arXiv.1901.09732>
- Tikk, D., Biro, G., Gedeon, T. D., Koczy, L. T., & Yang, J. D. (2002). Improvements and critique on Sugeno's and Yasukawa's qualitative modeling. *IEEE Transactions*

on *Fuzzy Systems*, 10(5), 596-606.

<https://doi.org/10.1109/TFUZZ.2002.803494>

Uahengo, F. D. L. (2014). *Estimating particle size of hydrocyclone underflow discharge using image analysis* [Master's thesis, Stellenbosch University]. SUNScholar.

https://scholar.sun.ac.za/bitstream/handle/10019.1/86365/uahengo_estimating_2014.pdf

van Deventer, J. S. J., Feng, D., Petersen, K. R. P., & Aldrich, C. (2003). Modelling of hydrocyclone performance based on spray profile analysis. *International Journal of Mineral Processing*, 70(1), 183-203.

[https://doi.org/10.1016/S0301-7516\(03\)00002-4](https://doi.org/10.1016/S0301-7516(03)00002-4)

van Loggenberg, S., van Schoor, G., Uren, K. R., & van der Merwe, A. F. (2016). Hydrocyclone cut-size estimation using artificial neural networks. *IFAC-PapersOnLine*, 49(7), 996-1001.

<https://doi.org/10.1016/j.ifacol.2016.07.332>

Wade, H. L. (2017). *Basic and advanced regulatory control: System design and application* (3rd ed.). International Society of Automation.

Wan, E. A., & van der Merwe, R. (2000). The unscented Kalman filter for nonlinear estimation. *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 153-158.

<https://doi.org/10.1109/ASSPCC.2000.882463>

Wang, Y., Velswamy, K., & Huang, B. (2017). A long-short term memory recurrent neural network based reinforcement learning controller for office heating

ventilation and air conditioning systems. *Processes*, 5(3), Article 46.

<https://doi.org/10.3390/pr5030046>

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8, 279-292.

<https://doi.org/10.1007/BF00992698>

Wei, D., & Craig, I. K. (2009). Grinding mill circuits — A survey of control and economic concerns. *International Journal of Mineral Processing*, 90, 56-66.

<https://doi.org/10.1016/j.minpro.2008.10.009>

Williams, R. A., Jia, X., West, R. M., Wang, M., Cullivan, J. C., Bond, J., Faulks, I., Dyakowski, T., Wang, S. J., Climpson, N., Kostuch, J. A., & Payton, D. (1999). Industrial monitoring of hydrocyclone operation using electrical resistance tomography. *Minerals Engineering*, 12(10), 1245-1252.

[https://doi.org/10.1016/S0892-6875\(99\)00109-0](https://doi.org/10.1016/S0892-6875(99)00109-0)

Wong, K. W., Ong, Y. S., Eren, H., & Fung, C. C. (2004). Hybrid fuzzy modelling using memetic algorithm for hydrocyclone control. *Proceedings of 2004 International Conference on Machine Learning and Cybernetics*, 7, 4188-4193.

<https://doi.org/10.1109/ICMLC.2004.1384574>

Xinhai Mining. (2016, March 17). *Hydrocyclone used for minerals (ores) classifying*,

Xinhai 2020 [Video]. YouTube.

<https://www.youtube.com/watch?v=OG4jy5XoBpo>

Yench, B., Metzner, G., & Lynch, A. (2015). Process control. In A. Lynch (Ed.), *Spectrum 21: Comminution handbook* (pp. 227-243). The Australasian Institute of Mining and Metallurgy.

Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 27, pp. 3320-3328). Curran Associates.
<https://proceedings.neurips.cc/paper/2014/file/375c71349b295fbe2dcdca9206f20a06-Paper.pdf>

Zhang, J., Tang, Z., Xie, Y., Ai, M., & Gui, W. (2020). Convolutional memory network-based flotation performance monitoring. *Minerals Engineering*, 151, Article 106332. <https://doi.org/10.1016/j.mineng.2020.106332>

Zhang, Y., & Liang, Y. (2016). Research on soft-sensor based on support vector regression for particle size of grinding and classification process. *2016 Chinese Control and Decision Conference (CCDC)*, 6708-6713.
<https://doi.org/10.1109/CCDC.2016.7532204>

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal Of The Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301-320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>

Every reasonable effort has been made to acknowledge the owners of copyright material. I would be pleased to hear from any copyright owner who has been omitted or incorrectly acknowledged.