Chemical Process Optimization Sample Bonsai Brain- AI Solution Spec

Project start/end dates: Click here to enter a date.

Authors: Denise Feirstien

Rows/questions highlighted in yellow are required before project can begin.

Customer	**Continuous Stirred Tank Reactor Optimization**	
Project Objective	*What task/process are you looking to improve using deep reinforcement learning?*	• A Continuous Stirred Tank Reactor (CSTR) is essentially a tank that has a stirring apparatus to continuously mix reactants inside. • There can be a number of inlets feeding the reactor, and an outlet removing the reacted product. • The reactor operates continuously, so during steady-state operation, we are continuously feeding it reactants and removing product with a fixed conversion rate. • This simple reactor is used in several applications, including chemical production and food and beverage manufacturing. • In this sample, we are looking at an exothermic reaction, that produces heat, meaning the reactor temperature must be controlled to prevent thermal runaway, when the reactor becomes uncontrollably hot. • The CSTR needs to operate under transient and steady state conditions. During continuous steady-state operation, the CSTR is producing a specified product. When the CSTR is starting up or transitioning to produce different output concentrations it is in transient state. • The transient state is difficult to control to reach the target concentration of the product while preventing thermal runaway. • **The objective of this Sample project is to train a brain(s) to provide coolant control settings for the CSTR to prevent thermal runaway during transition of residual concentration** **Figure 1**: CSTR Schematic 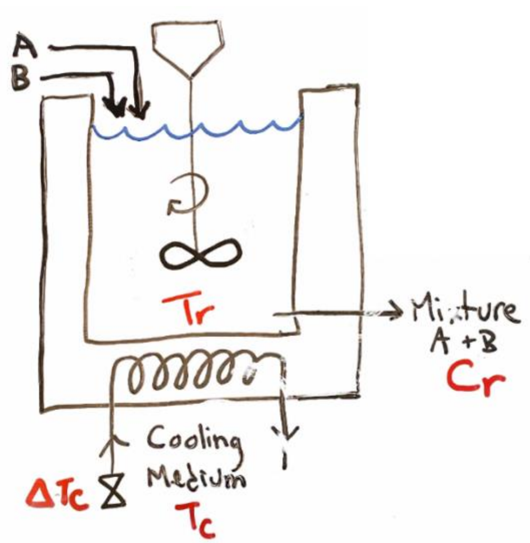

<table>
<tr>
<td>Business Value</td>
<td>What is the business value of improving the control/optimization of this system?</td>
<td>

Lower cooling costs
Prevent thermal runaway
Increase production by increasing conversion rate (decreasing residual concentration)

</td>
</tr>
<tr>
<td>Optimization Goal</td>
<td>What Key Performance Indicators (KPI) define the control or optimization of this system?</td>
<td>
<table>
<tr><th>Goal (KPI)</th><th>Units</th><th>Description</th></tr>
<tr><td>Minimize RMS error from reference Cr and Tr</td><td>None</td><td>Root mean squared (rms) error calculates the error from the reference at each timestep and aggregates the results throughout the chemical process (episode).
Minimize root mean square of error (Reference residual concentration – actual residual concentration)</td></tr>
<tr><td>avoid thermal runaway (>400K)</td><td>K</td><td>Avoid reactor temperature exceeding 400 K</td></tr>
</table>
</td>
</tr>
<tr>
<td>Current Methods</td>
<td>How do you currently control or optimize the system?</td>
<td>
<table>
<tr><th></th><th>Method</th><th>Level</th></tr>
<tr><td>☐</td><td>Human Operator</td><td></td></tr>
<tr><td>☐</td><td>Expert System</td><td></td></tr>
<tr><td>☒</td><td>Control Theory (PID, MPC)</td><td>Low-Level Control: Gain Scheduled PI, MPC</td></tr>
<tr><td>☐</td><td>Advanced Process Control (APC)</td><td></td></tr>
<tr><td>☐</td><td>Optimization Techniques</td><td></td></tr>
</table>
</td>
</tr>
</table>

Limitations of Current Methods

What are the challenges and limitations of the current method(s)?

Adaptive PID does not perform well when sensor noise is present

	Limitation	Description
☒	Ability to control well across scenarios / conditions	Extremely nonlinear process. Adaptive PI controller is only designed for one specific transition
☐	Multiple or changing optimization goals	
☐	Human Operator / Engineer Limitations	(see table below)
☒	Uncertainty in the measurement of the inputs or the process make it difficult to control or optimize.	Adaptive PI fails to prevent thermal runaway when >3% sensor noise exists
☐	Time to develop control or optimization system is prohibitive	

Human Operator / Engineer Limitations:

	Limitation	Details
☐	Difficulty managing many variables and dimensions.	
☐	Difficulty adapting to changing conditions	
☐	Large performance discrepancy between novice and expert operators	.
☐	Inconsistency across expert operators	

Machine Teaching Strategy

Heuristics

When the [environment variable list] trend in this direction or interact in this way,	This is what we think it means.	This is what you should do (to manipulate control actions).

Concept Network Decomposition

The first two heuristics above point to a decomposition that leverages two explainable strategies.

This architecture provides an additional layer of explain-ability so that the brain reports which strategy it is deploying in addition to the control actions.

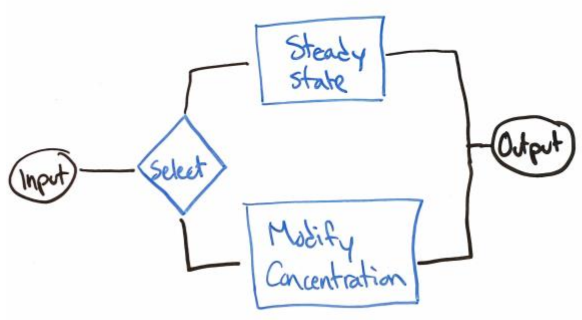

Concept 1: Steady State

- This concept learns to maintain reactor temperature while the reactor produces a constant residual concentration of 8.57 kmol/m3

Concept 2: Modify Concentration

- This concept learns to regulate the reactor temperature during transition

Selector Concept: Select Strategy

- This concept decides which strategy to deploy.

Control Actions

What actions will the brain need to output to control or optimize your system?

	Level	Number of Actions	Description
☐	Supervisory		The brain will provide supervisory set points.
☒	Low-level	1	Low-level control will remain with the APC controllers.

☐ **Low-Level Control System**: If the brain will provide supervisory control actions, is there a low-level system (APC, MPC, etc.) that must be included in the training loop? If yes, this control system must be integrated with the simulator and documented in the Simulation section below.

<table>
<tr>
<td></td>
<td></td>
<td>

Name	Data Type	Units	Control Frequency	Operating Range [min, max]	Description
dTc	Decimal	K/min	0.5 min	[-5, 5]	Change in coolant temperature at each timestep

☐ **Delayed Reward Scenario**: It takes a matter of minutes for most ore to move through the crusher and 15 minutes to get an average fragmentation time. We do not have a delayed reward scenario.

Rule of thumb: For optimal learning, system should settle to steady state within 1/10th of the control frequency.

</td>
</tr>
<tr>
<td>Constraints</td>
<td>What constraints are placed on the control actions by the system or the process?</td>
<td>

- The coolant temperature can change no more than 10 Kelvins/min, therefore dTc is constrained to [-5, 5] (for Ts=0.5min)
- The initial coolant temperature is 297.9798 K for an initial residual concentration of 8.57 kmol/m3

</td>
</tr>
</table>

Environment States	*What information do we need to pass to the brain about the system and its environment for the brain to learn to control or optimize the system?*	**Process Variables**

Name	Data Type	Source	Units	Measurement Frequency	Operating Range [min, max]	Description
Cr	Decimal	Simulation	Kmol/m3	[frequency]	[0.1, 12]	This is the actual residual concentration produced by the reactor
Tr	Decimal	Simulation	Kelvin	[frequency]	[10, 800]	This is the actual reactor Temperature
Cref	Decimal	Simulation	Kmol/m3	[frequency]	[0.1, 12]	The reference residual concentration is the desired concentration at every time step
Tref	Decimal	Simulation	Kelvin	[frequency]	[10, 800]	The reference reactor temperature is the desired reactor temperature at every time step
dTc	Decimal	Simulation	Kmol/m3	[frequency]	[-287.9798, 502.0202]	The change of coolant temperature from initial coolant temperature (at the beginning of simulation)
Tc	Decimal	Simulation	Kelvin	[frequency]	[10, 800]	The coolant temperature of the reactor cooling jacket. At each timestep, Tc = dTc + initial coolant temperature (297.9798 K)

Deep Reinforcement Learning	Deep Reinforcement Learning algorithms train agents to make sequential decisions which are assessed for the affect that each decision has on the environment.

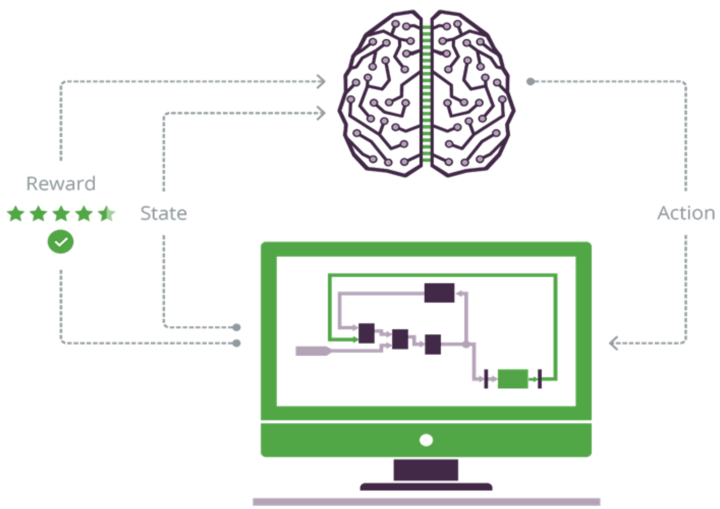

For each concept that we will train using Deep Reinforcement Learning, we outline the sequential decision

Concept	Action	State: How does the Environment change when the control actions are taken?	Reward	Configuration: What do we need to vary in the training to ensure that the brain works well across scenarios?
Steady State	Change coolant temperature	Each time a decision is made to the coolant temperature, the Cr and Tr change.	• Minimize error of Cr • Avoid thermal runaway	• Noise
Modify Concentration	Change coolant temperature	Each time a decision is made to the coolant temperature, the Cr and Tr change.	• Minimize error of Cr • Avoid thermal runaway	• Noise
Selector	Change coolant temperature	Each time a decision is made to the coolant temperature, the Cr and Tr change.	• Minimize error of Cr • Avoid thermal runaway	• Noise • When transition from Steady State to Transient conditions occurs

<table>
<tr>
<td>Configuration Scenarios</td>
<td>What scenarios should the trained brain be able to control across?</td>
<td>
Deep Reinforcement Learning (DRL) can produce brain(s) that control well across a wide range of scenarios and is particularly suitable for situations where the distribution of the variables in the configuration scenarios is unknown and / or non-linear.
<table>
<tr><th>Configuration Variable</th><th>Range [min, max]</th><th>Description</th></tr>
<tr><td>Noise</td><td>[0, 100]</td><td>Amount of gaussian noise added to the Cr and Tr of the system</td></tr>
<tr><td>Cref signal</td><td>[1, 5]</td><td>The Cref signal defines the desired trajectory of the Cref and Tref for an episode. The following 5 singals are defined:
1: Concentration transition --> 8.57 to 2.000 over [0, 0, 26, 90] (minutes) – 0 delay
2: Concentration transition --> 8.57 to 2.000 over [0, 10, 36, 90] (minutes) – 10 min delay
3: Concentration transition --> 8.57 to 2.000 over [0, 20, 46, 90] (minutes) – 20 min delay
4: Concentration transition --> 8.57 to 2.000 over [0, 30, 56, 90] (minutes) – 30 min delay
5: Steady State --> 8.57</td></tr>
</table>
Training Episode Length: 180 control actions (90 minutes simulated)

Benchmark Episode Length: 180 control actions (90 minutes simulated)
</td>
</tr>
<tr>
<td>Success Criteria</td>
<td>What criteria will we use to determine the success of the project and how will we measure that success criteria?</td>
<td>
<table>
<tr><td>KPI</td><td>Error of Residual Concentration (RMS of error from reference), Thermal Runaway condition</td></tr>
<tr><td>Benchmark Comparison</td><td>The brain will be compared to a gain-scheduled PI controller</td></tr>
<tr><td>Benchmark Scenarios</td><td>
<table>
<tr><th>Configuration Variable</th><th>Units</th><th>Priority</th><th>Range or Description</th></tr>
<tr><td>Noise</td><td>%</td><td>1</td><td>0%, 5%, and 10% noise</td></tr>
<tr><td>Reference Single</td><td>none</td><td>1</td><td>Benchmark will only be run for Cref single 2</td></tr>
</table>
</td></tr>
</table>
</td>
</tr>
</table>

		Benchmark Procedure |	Procedure	Duration
☒	Simulation	[Benchmark Duration in Simulation]
☐	A/B Testing on Live System	[A/B Testing on Live System] | **Optimization Improvement**: [success criteria expressed in % improvement over current methods] **Return on Investment (ROI)**: [success criteria expressed in return on investment (ROI)] **Project Readout and Deliverables**: [expected deliverables besides the brain(s) and a PowerPoint readout report] ☒ Is the benchmark controller or performance noted above currently ready?
Simulation		See tables below

Benchmark Procedure

	Procedure	Duration
☒	Simulation	[Benchmark Duration in Simulation]
☐	A/B Testing on Live System	[A/B Testing on Live System]

Optimization Improvement	[success criteria expressed in % improvement over current methods]
Return on Investment (ROI)	[success criteria expressed in return on investment (ROI)]
Project Readout and Deliverables	[expected deliverables besides the brain(s) and a PowerPoint readout report]

☒ Is the benchmark controller or performance noted above currently ready?

Simulation

Readiness

Delivery Date	[Sim Delivery Date]
Validation Date	[Sim Validation Date]
Sim Builder	[Sim Builder]
Integration with Microsoft Machine Teaching Service	SDK3

Type

Vendor	
Product (Version)	N/A, custom sim
Software Language API Interface	Python
Speed	Less than 1 second per iteration

<table>
<tr><td></td><td></td><td>Modeling Method</td><td>
<table>
<tr><th></th><th>Method</th><th>Description</th></tr>
<tr><td>☒</td><td>Physics Based (First Principles)</td><td>Mathwork Simulink model of CSTR</td></tr>
<tr><td>☐</td><td>Discrete Event</td><td></td></tr>
<tr><td>☐</td><td>Surrogate Model</td><td></td></tr>
<tr><td>☐</td><td>Model from Data</td><td>The amount (number of rows) of data required to create a simulation model from data varies, but use the following rule of thumb as an absolute minimum: the number of possible states x the number of possible actions. For example, if there are 10 possible actions and 100 possible states, you’d need 1,000 rows of data at minimum to build a model.

Model Accuracy & Robustness
The model should be validated across the ranges for each of the control actions and environment states listed above. Enter the accuracy of the model for each of the features.
<table>
<tr><th>Feature</th><th>Accuracy</th></tr>
<tr><td>[One row for each control action and environment state]</td><td>[% Error]</td></tr>
</table>

State Space Completion
<table>
<tr><th></th><th>Procedure</th><th>Rows of Data</th></tr>
<tr><td>☐</td><td>State Space Parameter Sweep</td><td>[data volume]</td></tr>
<tr><td>☐</td><td>Synthetic State Space Estimation</td><td>[data volume]</td></tr>
</table>
</td></tr>
</table>
</td><td></td></tr>
<tr><td></td><td></td><td>Connection</td><td>Can we exchange messages (input and output) with the simulation model at the simulated control frequency?

Is a high-level control system diagram from sensors to actuators available?
Are there any other pieces required, beside the simulator, to run the training loop?

[Can this software connect for input and output on the inner loop?]</td><td></td></tr>
<tr><td></td><td></td><td>Configuration</td><td>Can we input the configuration scenarios programmatically into the simulation model?

[Can we input configuration scenarios programmatically into the simulation model?]</td><td></td></tr>
<tr><td></td><td></td><td>Parallelization (Licensing)</td><td>Can we run 10, 100 or 1000 copies of your simulation in the Azure cloud?

[Can we run the simulation in the Azure Cloud?]</td><td></td></tr>
</table>

		Simulation to Reality — *Has the simulation model been used to design a control system, an optimization system or used by human operators in production to control the system.* No. *What is the error percentage that describes the accuracy between the simulation model and the real system across all scenarios and equipment that will be controlled by the brain?* [% error] *Do you plan any simulator upgrades, especially if it will need to be upgraded for use with Bonsai?* **Simulation Validation** — *Are there any major assumptions in the sim that would change the sim dynamics as compared to the real-world dynamics?* *Can you provide the validation data against your sim?* *Are there special, external libraries?* *Does the model connect to external data sources?* *How many workarounds are needed to setup the model to run headlessly, ideally using only parameters on the top level agent (Main)?*
Supplementary Decision Models	Will Machine Learning (ML) models or other decision-making technology be used to supplement the environment state from the simulator?	(see table below)

Type	Training Data	Model Accuracy	Description
Liner Predictor	[training data]	[model accuracy]	[description]
ROM Bin Input Predictor	[training data]	[model accuracy]	[description]
COB Bin Output Predictor	[training data]	[model accuracy]	[description]
Rock Hardness Fragmentation Predictor	[training data]	[model accuracy]	[description]
Maximum Safe Crusher Gap Predictor	[training data]	[model accuracy]	[description]

Delivery Date	[Delivery Date]
Validation Date	[Validation Date]
Model Builder	
Integration with Microsoft Machine Teaching Service	Python SDK2

Deployment	*How will the brain interface with your system? (select & respond to one or multiple options below)*	Is the deployment interface and protocol defined and ready? If it does not exist, what is the delivery date?

☐	**Decision Support**	*Human engineers, operators or analysts will continue to control and automate my system augmented by brain decisions.* ☐ **Cloud Deployment** ☐ **Edge Deployment** ☐ **Embedded Deployment** Integration with OT environment will be through existing IoT/OPC Gateway Edge infrastructure enabled to host containers, within the IT environment.
☐	**Direct Control**	*The brain will connect to the system directly to automate the control or optimization.* ☐ **Cloud Deployment** ☐ **Edge Deployment** ☐ **Embedded Deployment** For Edge and Embedded Deployments:

	Decision Delivery Mechanism	Description
☒	Decision Support UI	
☐	Spreadsheet or other mechanism	
☐	Integration with current reporting system	

Device Type	[Device Type]
Number of Devices	[Number of Devices]
Device Lifecycle	[Device Lifecycle]
Docker Support	[Docker Support]
Processor	[Processor]
Connection Protocol	[Connection Protocol]
Integrator	[Integrator]
Integration Delivery Date	[Integration Delivery Date]

<table>
<tr><td>Team</td><td></td><td>
<table>
<tr><td>Executive Sponsor</td><td>[Executive Sponsor Name]</td></tr>
<tr><td>Machine Teacher</td><td>[Machine Teacher Name]</td></tr>
<tr><td>Data Scientist (Optional)</td><td>[Data Scientist Name]</td></tr>
<tr><td>Subject Matter Expert</td><td>[Subject Matter Expert Name]</td></tr>
<tr><td>Simulation Expert</td><td>[Simulation Expert Name]</td></tr>
<tr><td>Deployment Expert</td><td>[Deployment Expert Team]</td></tr>
<tr><td>IT</td><td>[IT Contact Name]</td></tr>
<tr><td>Project Team</td><td>[]</td></tr>
<tr><td>Services Partner</td><td>[Services Partner Team]</td></tr>
<tr><td>Microsoft Applied AI Engineer</td><td>[Project Applied AI Engineer]</td></tr>
<tr><td>Microsoft Technical Program Manager</td><td>[Project Technical Program Manager]</td></tr>
<tr><td>Microsoft Account Team</td><td>[Account Executive, Account Technical Strategist Names]</td></tr>
<tr><td>Microsoft CSA</td><td>[CSA Name]</td></tr>
</table>
</td></tr>
<tr><td>Azure Infrastructure</td><td></td><td>
<table>
<tr><td>Azure Subscription</td><td>[Subscription ID]</td></tr>
<tr><td>Other Azure Services Required</td><td>[List of Azure Services]</td></tr>
</table>
</td></tr>
</table>

		Will Microsoft be given access to the customer's Azure subscription?