# GOOML - Finding Optimization Opportunities for Geothermal Operations

## Preprint

Paul Siratovich,[1] Grant Buster,[2] Nicole Taverna,[2] Andrea Blair,[1] Jon Weers,[2] Michael Rossol,[2] and Jay Huggins[2]

*1 Upflow Limited, New Zealand*
*2 National Renewable Energy Laboratory*

*Presented at the Stanford Geothermal Workshop (SGW)*
*Stanford, California*
*February 7-9, 2022*

# GOOML - Finding Optimization Opportunities for Geothermal Operations

## Preprint

Paul Siratovich,[1] Grant Buster,[2] Nicole Taverna,[2] Andrea Blair,[1] Jon Weers,[2] Michael Rossol,[2] and Jay Huggins[2]

**NOTICE**

# GOOML - Finding Optimization Opportunities for Geothermal Operations

**Paul Siratovich[1], Grant Buster[2], Nicole Taverna[2], Andrea Blair[1], Jon Weers[2], Michael Rossol[2], Jay Huggins[2]**

**[1]Upflow Limited, New Zealand**

**[2]National Renewable Energy Laboratory (NREL)**

## Keywords

*Machine Learning, GOOML, Optimization, Operations, Forecast, Steamfield, Geothermal, Digital Twin*

## ABSTRACT

Geothermal Operational Optimization with Machine Learning (GOOML) is a transferable and extensible component-based geothermal asset modeling framework that considers complex steamfield relationships and identifies optimization prospects using a data-driven approach. We have used this framework to develop digital twins that provide steamfield operators with an operational environment to analyze and understand historical and forecasted power production, explore new steamfield configuration possibilities, and seek optimal asset management for real world applications.

The GOOML modeling software is built on a generic component-based systems framework that allows for both historical and forecast analysis. A GOOML model can perform historical data-assimilation using first-principal thermodynamics to create a meaningful data model. Historical production data can then be coupled with a forecast framework to train machine-learning models of steamfield components to predict future outputs. This modeling environment enables digital exploration of steamfield design configurations and operational scenarios.

GOOML digital twins have been developed for steamfields in New Zealand and the United States representing differing power generation and field conditions. These digital twins have been validated by comparing hindcast predictions against historical production data. Reinforcement learning experiments were conducted to demonstrate the ability to programmatically explore the operations space using machine learning agents. Our initial results are compelling; two to five percent increases in annual energy production were demonstrated by the GOOML models with no additional infrastructure build required.

GOOML offers a new approach to geothermal operations by applying state-of-the-art machine learning algorithms, comprehensive data analytics, and interaction with digital twins. Through application of these tools, operators will realize greater availability and higher net generation which will increase the cost effectiveness of geothermal energy projects.

1

## 1. Introduction

The development of geothermal fields for power production is a capital and resource intensive endeavor. Drilling, construction, and commissioning are all focused on bringing plants online and delivering power to the grid. Once this initial phase has passed asset issues begin to emerge. The surface infrastructure may need additional changes or redesign and additional wells may be needed to fuel stations. As a result, aged geothermal operations may be run at conditions that are far from the original design constraints. Operators are constantly searching for ways to optimize their infrastructure assets while minimizing additional capital or operational expenditure. For this reason, operators will look to their datasets for trends, inefficiencies, and opportunities to enact changes that may be beneficial to the operation of their fields.

Geothermal power projects must maintain development flexibility to manage production and environmental requirements and limits over time, and due to their long-term nature, surface infrastructure becomes more complicated with an increasing mix of new and old technologies, maintenance needs, changes to management regimes and commercial priorities, as well as socio-political contexts. Thus, managing data and information from geothermal assets, in a systematic and holistic way has always proven to be a challenge for operators.

This need to interrogate and understand complex data from a variety of sources is what has yielded the GOOML project. By using a data-science approach coupled with significant knowledge of geothermal operations, an opportunity exists for a new tool to be developed that operators can use to test scenarios and challenge current paradigms. Through a multi-disciplinary approach, we can provide insight and scrutiny of data to minimize disruption from maintenance, both planned and unplanned, and provide a comprehensive view of the operation through interrogation of data at a single source. GOOML is a digital environment that we have developed that allows the interrogation and manipulation of geothermal operational data with the aim to minimize down-time, identify anomalous events and provide a scenario-analysis tool all in a digital environment.
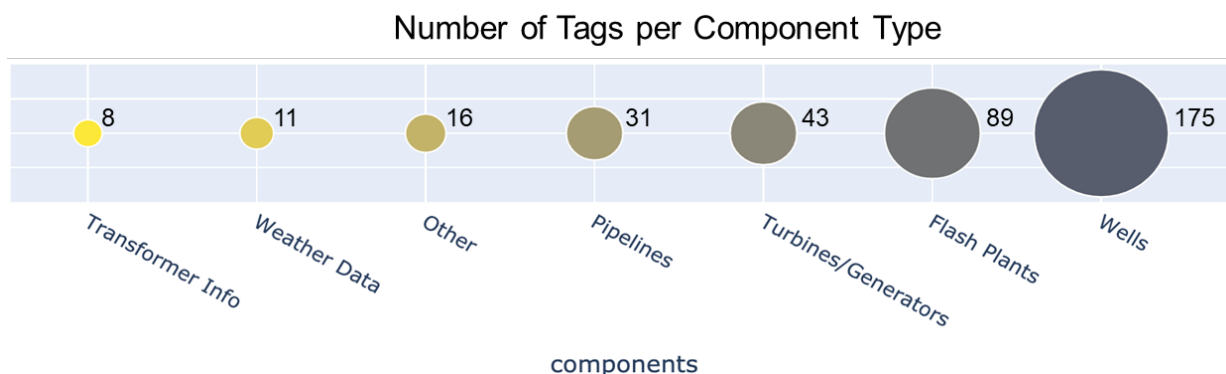
This platform will enable geothermal operators to interrogate their historical performance and compare to optimal conditions to find gaps or target areas for optimization. Further, through the project we have developed several tools that could be used in building hypothetical scenarios, help guide integrated steamfield design and provide a platform to test for multiple operational scenarios. We will discuss the data sources, model architecture and machine learning approach, initial results of the digital steam field twins, model exploration, before finally providing conclusions and suggested future work.

## 2. Data Sources

We obtained data sourced from real-world geothermal operations from Contact Energy's (Contact) Wairakei Field, Ngati Tuwharetoa Geothermal Assets' (NTGA) Kawerau Field and Ormat Technologies' (Ormat) McGinness Hills Field. These data span multiple years and cover nearly all relevant operating conditions for modern geothermal fields. The timespans and operating conditions covered by these data, as well as the complexity of the steamfields selected, were critical to the development of GOOML. Comprehensive datasets incorporating as many real-world operating conditions and steamfield configurations as possible were necessary to properly inform machine learning (ML) experiments and reduce potential training biases. Datapoints are sourced from several steamfield components, including but not limited to wells, pipelines, flash plants

(separators), and turbine-generators (T/Gs). The breakdown of this data by component type is shown in Figure 1.

To properly build a digital twin using the GOOML framework the operational data was curated into a consistent time-series format. Data curation was a recursive process that consisted of the following steps: 1) acquiring data from power plant operators, 2) digesting data to get a basic understanding of what is included, 3) transforming data into a machine-readable format so that it may be visualized, quality checked, and eventually used in the GOOML models, 4) simple data cleaning to resolve issues that may not be solved with additional data, 5) identifying significant data gaps and apparent anomalies, and 6) discussion with modelers to identify additional data needs and to resolve any data gaps or anomalies.

### Number of Tags per Component Type



| 8 | 11 | 16 | 31 | 43 | 89 | 175 |

Transformer Info | Weather Data | Other | Pipelines | Turbines/Generators | Flash Plants | Wells

components

**Figure 1: Initial snapshot of the type of data used in development of GOOML models with a heavy focus on wells and steam-water separators (flash plants).**

## 3. Model Architecture and Machine Learning Approaches

GOOML is a modular, component-based model that allows for the creation of a digital twin for any steamfield through the configuration of easily customizable digital components that function as extensible building blocks for replicating real-world steamfield components such as flash plants, join junctions, and generators. These generic building blocks are encoded with the basic data streams required to describe their respective components. For example, a flash plant component is defined by its two-phase input stream, its separated output flows and enthalpies, and its physical design and dimensions. Best practice object-oriented software design is used to create abstract building blocks on which historical data assimilation models or future forecast regression models can be built. Continuing the example, our flash plant forecast model is a machine-learning based regression model that is based on an abstract flash plant component template, which is in turn based on a generic split junction component. Various component models can then be connected into a component-based systems model as shown in Figure 2. This generic component-based systems framework is what makes GOOML a powerfully generic and extensible tool.
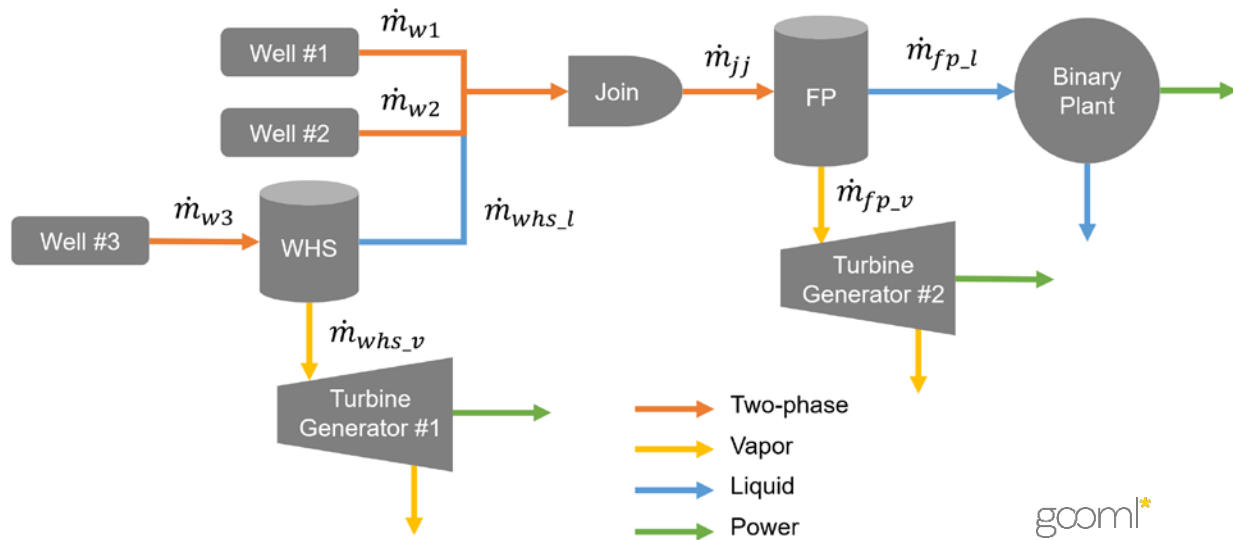
3

**Figure 2: An overview of how the GOOML component models are connected within a simplified steamfield system network.**

### 3.1 Historical Data Model

The historical data model is a data assimilation framework built on the generic GOOML component-based systems modeling ecosystem. Using this model is the first step an analyst would take to build a GOOML model for a steamfield.

First, a system configuration is developed to organize the general topology of the steamfield network as shown in Figure 2. Input time-series data is automatically assigned to each component object within the system (e.g., a well gets assigned its respective pressure, mass flow, and enthalpy data streams). The historical data model attempts to enforce conservation of mass and energy equations by manipulating two-phase flow estimates while preserving trusted and accurate single-phase flow measurements. The model also fills in any "gaps" using first principal thermodynamics. For example, the two-phase properties at the inlet of a flash plant are often not directly measured. The historical data model automatically calculates these properties based on the upstream mass and energy flow coupled with a calculation of the pressure drop across the flash plant. The result is a cleaned and standardized dataset in a relational structure from which algorithms can be developed and trained.

### 3.2 Regression Modelling Approaches for Forecasting and Hindcasting

Using the historical data model as a foundation, regression models can be trained to predict how each component interacts with steam, liquid, or two-phase flow. A forecast system model can utilize several regressions specific to individual steamfield components or use a generic model that represents a whole class of components. Once the models are trained, the forecast system can be run to predict future operations or to hindcast historic operations for validation purposes. It is important to note that the GOOML forecast system cannot predict every parameter, and basic inputs are still needed, such as estimated well head pressures.

4

The GOOML forecast system typically starts with an estimate of the well head pressure that is used in conjunction with Tracer Flow Testing (TFT) equations to estimate a well's two-phase mass and energy output. From there, the forecast system operates as a simple feed-forward network. The fluid flow is propagated throughout the system, and downstream component-based regression models perform thermodynamic operations on the fluid. Eventually, the fluid flow reaches the last components in the network, typically some sort of power-generating unit or steam end-user. The regression models for these components yield an estimate of power or steam production that serves as a metric for overall steamfield performance. Examples of this can be seen in subsequent sections. This method for calculating a forecast using this generic systems approach allows exploration of different component interactions, uncertainty ranges and hypothetical scenarios.

3.2.1 Component-Specific Regression Modelling Approaches

We have used several different approaches to build regression models for the respective steamfield components. In many cases, the same approach was not suitable for each component and so the methods used are elaborated below.

**Wells** are forecasted using linear or exponential mass and enthalpy decline trends based on historical TFT measurements. Well head pressure is typically estimated from recent historical data during nominal operations or can be taken directly from the historical record for a pure validation of the other regression models. Using these field-derived datasets, the regressions are trained to give an accurate representation of how performance from wells decline over time within the forecast model.

**Flash plants** (steam/water separators) are interesting components that utilize phase change and mechanical separation to convert a two-phase fluid flow into a high-quality saturated steam and saturated water outputs. The complex physical processes happening in a flash plant are difficult to model using simple regression methods. As a result, we found Deep Neural Networks (DNNs) to be a useful tool to model flash plants.

The model architecture for the flash plants has three hidden layers each with 128 nodes. The layers are fully connected but trained with 50% dropout. All nodes use the rectified linear unit (RELU) activation which was chosen based on our attempt to linearize some of the training features. We originally performed a full gaussian hyperparameter search to optimize the model architecture but decided against the "optimized" model architecture because it resulted in saliency maps that were completely unexplainable, non-intuitive, and likely non-physical. This "final" architecture of 5x128 with 50% dropout resulted in similar validation error compared to the "optimized" model but also produced explainable and highly intuitive saliency maps that we deemed to be more likely representations of the physical phenomena.

The inputs to this model component are mass flow, enthalpy, pressure, theoretical pressure drop, theoretical flash fraction, the cyclone design number, and the input velocity. The DNN predicts a non-dimensional separation efficiency based on these inputs. For this work, we define the flash plant non-dimensional separation efficiency as the steam output mass flow divided by the total input mass flow. This non-dimensional number ranges from zero to one and can be used as a direct multiplier on the input flow to calculate the steam output flow.

5

**Turbine Generators**: Solving for T/G relationships proved to be an initially challenging problem as simple steam-flow-to-power relations (e.g. "Willans curves", see Church 1928) typically used in turbine design did not accurately reflect the steam to power output relationships. We also attempted to fit a sigmoid relationship with little success. For GOOML, a multi-linear regression with input features being mass flow, enthalpy flow, heat sink temperature, and temperature differential yielded the most accurate fit to historical conditions.

**Binary units**, like the T/G relationships proved not to be as simple as a mass flow and output relationship. To fit the historical model, we had to integrate another multi-linear regression that incorporated: enthalpy flow, temperature differential and mass flow contribution fractions from all upstream flash plants. The resultant relationships provided relatively good model accuracy to forecast the historical conditions.

3.2.2 Model Assumptions

It is important to discuss the assumptions that are made by simplifying a complex steamfield to the component-based system network shown in Figure 2. Because of the magnitude of the engineering complexity in a large steamfield, it is impractical (and likely nearly impossible) to get reliable data for the full system. Every pipe is not fully instrumented for all thermodynamic variables and every pipe does not have readily available design data (e.g. isometric drawings, dimensions, and hydraulic loss information).

For example, by using basic thermodynamics, we have attempted to solve for missing thermodynamic variables that would ideally be measured directly. Some data gaps inevitably exist in operational data (e.g. sensor failure, data loss, erroneous readings, etc.) and we have used empirical solutions to fill these gaps; the results using this approach are encouraging and are what we would consider as complete as possible.

Pressure is an important parameter that is often not measured throughout the full steamfield. For example, when working with a forecast model, the well head pressure is often simply assumed to be a nominal operating discharge pressure, unless the well is fitted with a control valve. Immediately downstream, calculating two-phase pressure drop between well and a flash-plant inlet would require significant information on both the quality of the two-phase fluid and the physical design of the piping system. Even then, calculations of two-phase pressure drops are often estimations at best. Instead, we have made the simplifying assumption that operators will be able to maintain a setpoint pressure at the turbine inlet. That pressure should be close to the steam output pressure at the flash plant (based on estimates of single-phase flow pressure drop from the flash plant to the turbines). We then iteratively solve for the pressure drop across the flash plant to get the inlet pressure. Combined with the enthalpy estimates delivered by the wells, this fully defines the thermodynamic state at the flash plant inlet.

## 4. Initial Results by Steam Field Component

The results presented here use the GOOML forecast model with historical well pressures, turbine pressures, and operator actions to validate the accuracy of the forecast models against historical data. This practice of "hindcast" validation shows how accurate the GOOML models can be with nominal data inputs.

Through an iterative process, we have been able to achieve a close match with our forecasts to that of the actual field performance. Overall, we can match system outputs in the model with relatively good agreement with the original datasets; c.a. 6.5% mean absolute error (MAE) for a two-year mass-take measurement. Table 1 elaborates the MAE and the mean bias error (MBE) for the significant field components of the Wairakei Steamfield.

It should be noted that the GOOML forecast model can accept bias correction terms to better match a trained regression model to historical data. However, in the interest of a truthful evaluation of model performance, no manual bias correction was performed here.

| Component | MAE | MAE (%) | MBE | MBE (%) |
|---|---|---|---|---|
| *Total Mass-Take* | 774 tonnes/hour | 7.95 | -367 tonnes/hour | -3.77 |
| *Total Separated Steam Flow* | 130 tonnes/hour | 5.11 | -79.73 tonnes/hour | -3.11 |
| *Steam Turbine Power Generation* | 17.16 MWe | 5.35 | -13.98 MWe | -4.36 |
| *Binary Generation* | 1.29 MWe | 11.55 | 1.11 MWe | 9.93 |

**Table 1: Mean Error Measurements of the Forecast Model for the Wairakei Steamfield Major Components over a 3-year period. Negative bias errors imply that the GOOML forecast model is under-predicting.**

### 4.1 Wells

Production wells comprise the most upstream portion of the GOOML environment. Wells can be either single phase liquid, two-phase or dry steam wells. The wells are forecasted using linear or exponential mass and enthalpy decline trends based on historical TFT measurements. Well head pressure is typically estimated from recent historical data during nominal operations. (Figure 3).
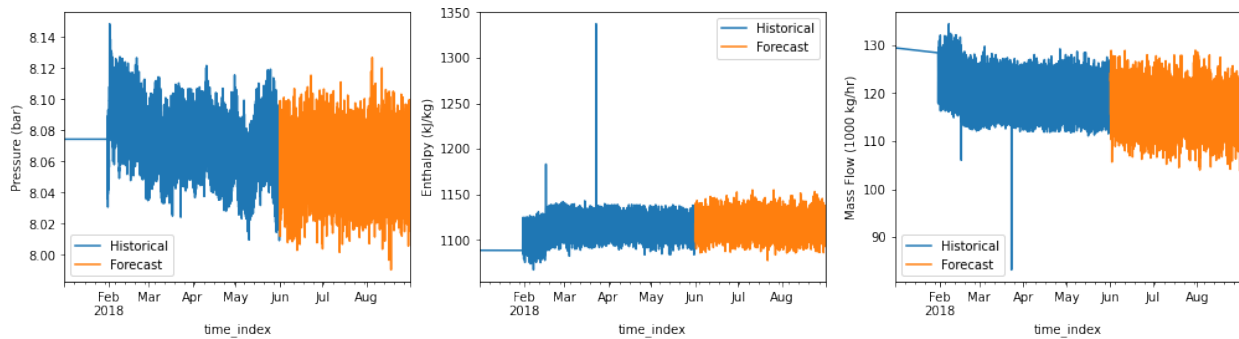


**Figure 3: Forecast matching from the historical model operational data for a well and the future forecast decline.**

By enforcing conservation of mass and enthalpy in the system, we have been able to create models of wells that match historical operations. This tool allows us to model well outputs and fluid flows to downstream components in the forecast space. This is seen in Figure 4 where a total field mass extraction is shown. There is relatively good agreement between actual (historical model) and forecast (forecast mean) models to represent the geothermal operation of the wells we have interrogated. By achieving good historical matching from the wells, we can then look to optimize components further downstream, specifically flash plants and generators.
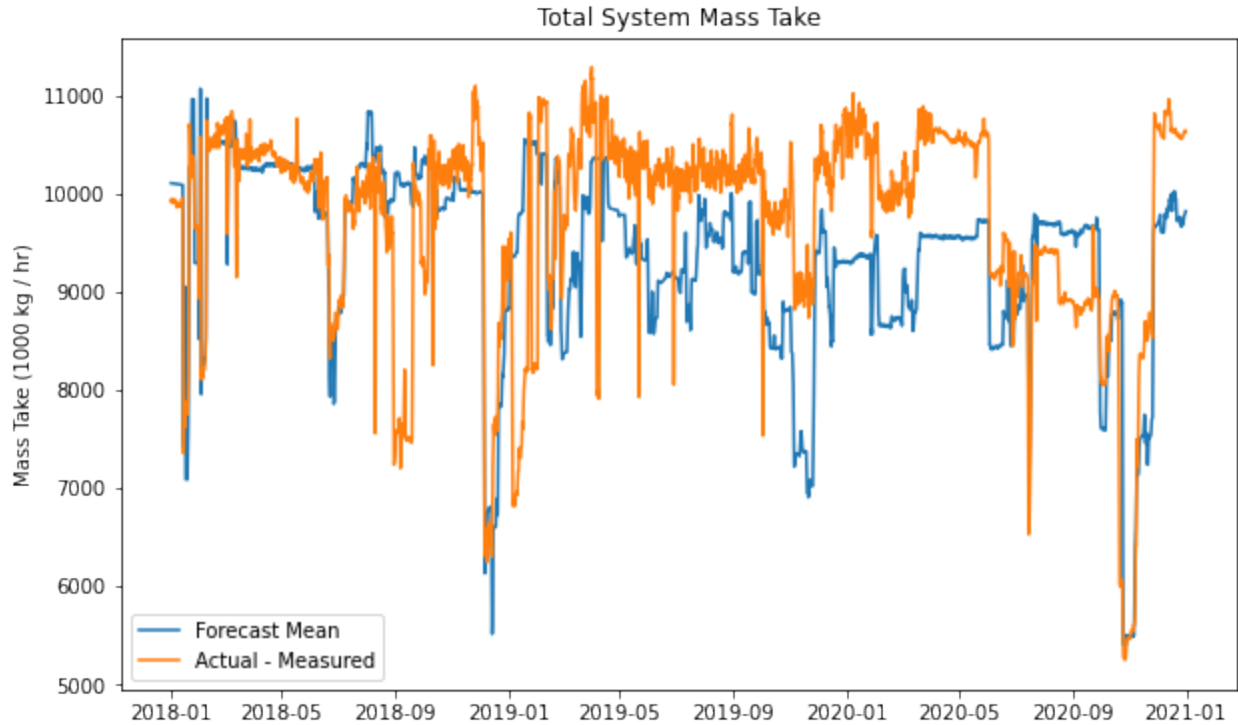
**Figure 4: Results of mass take from 3 years of historical data coupled with the forecast model mean results. Overall a good agreement with the historical model shows a mean absolute error (MAE) of 7.9%.**
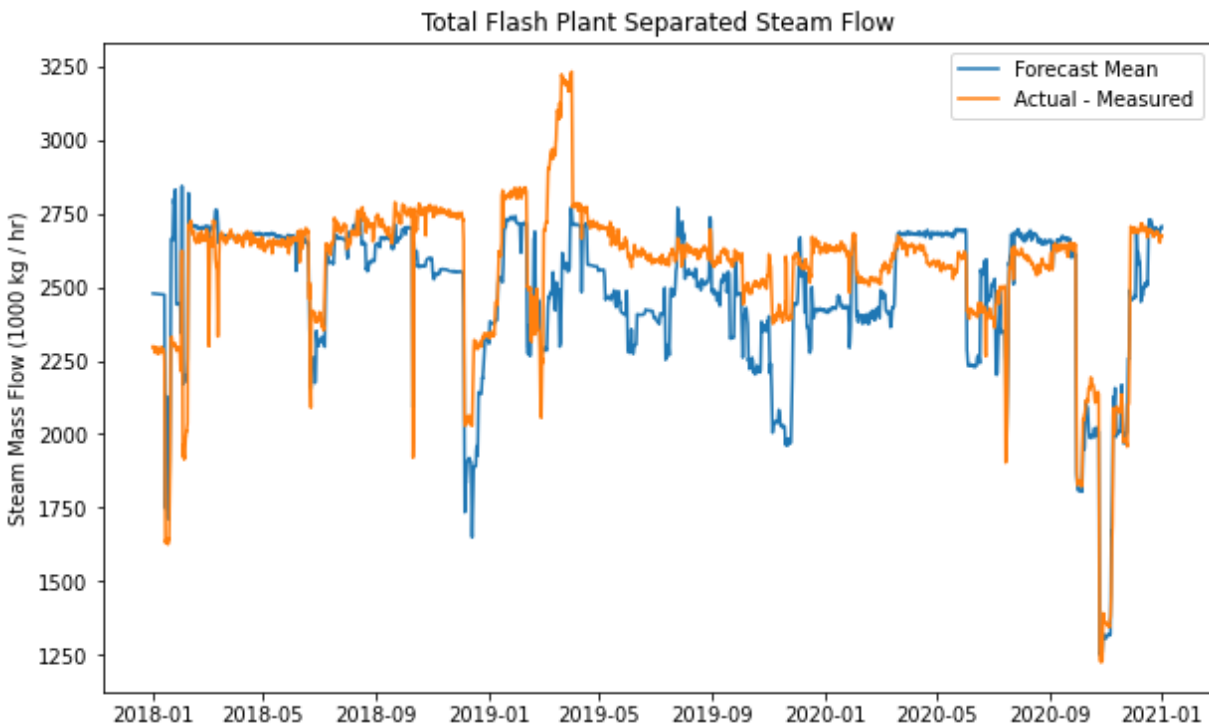
### 4.2 Flash Plants

Deconvoluting a geothermal flash plant process required significant effort to develop models that can accurately represent the conditions seen at the Wairakei and Kawerau Fields. As described in Section 3.2.1, a deep neural network was used to predict each flash plant's separation efficiency. This in turn can be used to compute the output steam flow from the flash plant based on the input values.

Along with accurately predicting flash plant operation, (discussed below), another advantage of using a deep neural network to model the flash plants is the ability to model the full operational range of each flash plant. This allows for the visualization of flash plant separation efficiency based on massflow and enthalpy, indicated in the multi-dimensional plots in Figure 5. The resulting prediciton surface then yields a separation efficiency that operators can target to attempt to achieve an optimal output from each flash plant.

8
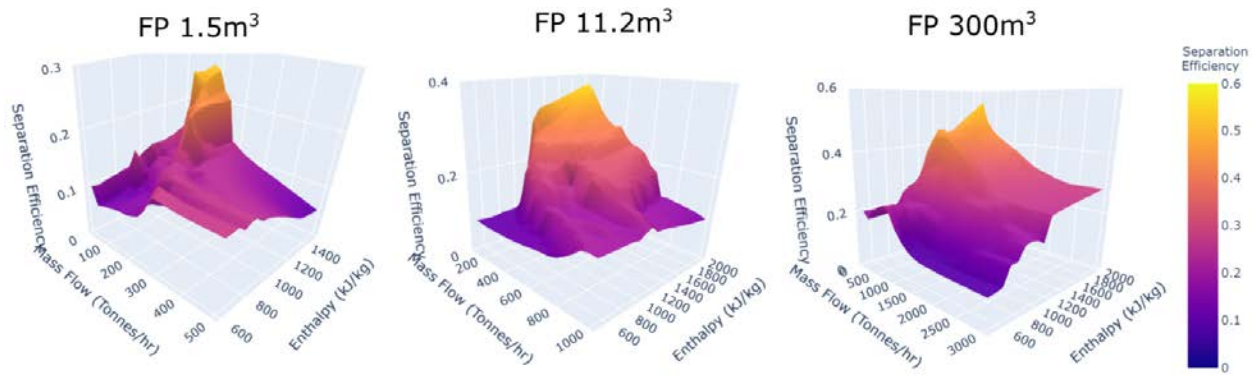
**Figure 5: Flash plant prediction surfaces for three various sized flash plants developed in the GOOML environment. Each plant has a unique optimal separation efficiency driven by upstream thermodynamic conditions and the mechanical design of the flash plant itself.**

For the Wairakei and Kawerau steamfields, the matching of flash plant outputs can be benchmarked by the steam-flow in pipelines downstream of the flash plants themselves. By comparing our forecast efficiency and thus, steam output, we can assess how well we have matched the true operating conditions of the plants. Figure 6 shows the error and matching of the historical operating conditons and the predicted outputs from our model. Using these outputs, operators may be able to predict where operational efficiency could be gained and finding the ideal conditions based on changes within a field.



**Figure 6: Flash plant historical vs. hindcast results showing a good match for stable periods of operation.**

9

## 4.3 Power Generation

One of the aims of the GOOML project is to provide operators interrogation of field configurations that would lead to optimal power output from their systems. Through accurate and interconnected modeling of wells, pipelines, and separators, we were able to create a forecast model for system-wide power generation. For this forecast, we modeled nine separate steam-fed turbine-generator units and two binary bottoming units. The results of the historical model and the hindcast are shown in Figure 7.

This combined power-output modeling gives a tangible measure of how accurate the GOOML model is in representing the full production system. For the full system power, we have achieved an MAE of 5.3% across three years of production history. We do see that the GOOML hindcast model does somewhat underpredicts the power generation compared to the historical model but are comfortable that we have achieved a good match. We expect that the systematic bias error is due to operator action that is not well captured in the historical numerical data. While this could be easily corrected in the GOOML modeling framework through the application of a bias correction factor, Figure 7 does not include such corrections to demonstrate a truthful evaluation of the trained regression models.
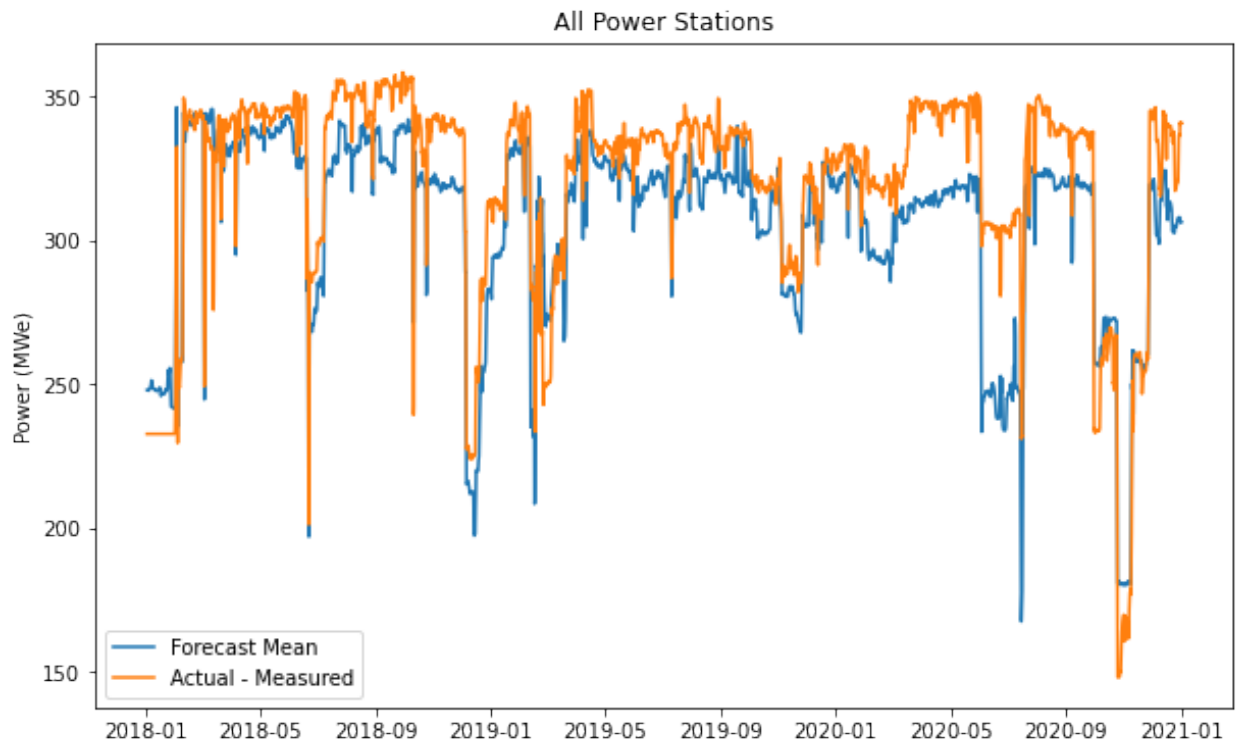


**Figure 7: Forecast and historical matching of power production for nine generation units. Overall good agreement is achieved within the GOOML environment.**

10

## 5. Model Exploration

Once we established that our model represents real-world conditions from initial modeling, we wanted to explore what could be further developed using this framework. Additionally, we needed to ensure that the model is extensible to other fields and conditions and can be used to test various scenarios.
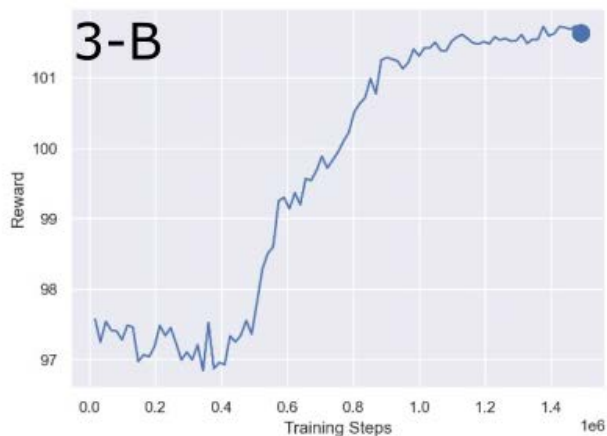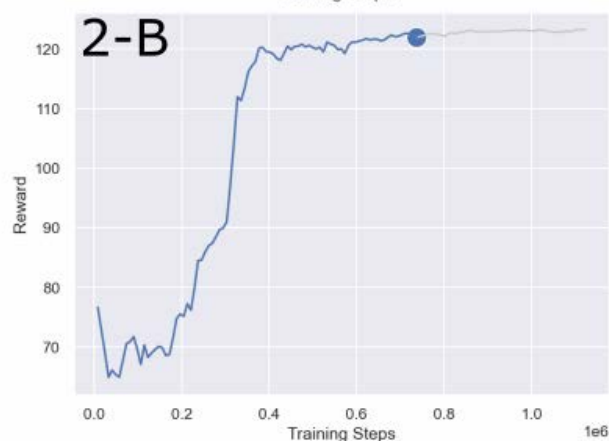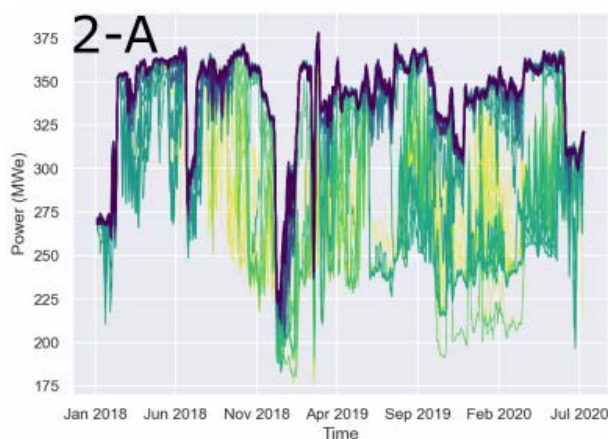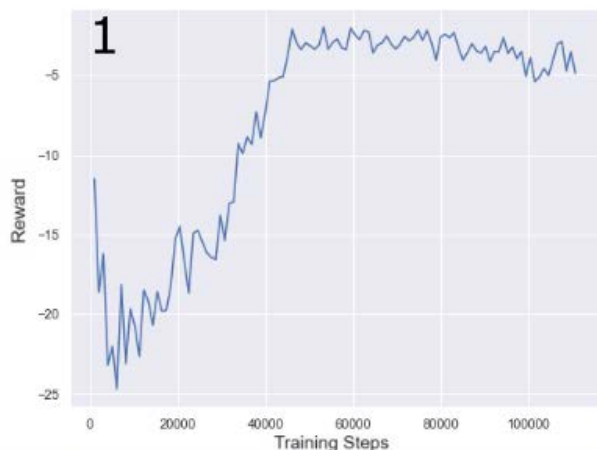
### 5.1 Reinforcement Learning

Artificial intelligence via reinforcement learning was applied to the GOOML environment to explore if an autonomous decision-making tool might provide insight to geothermal operations and test long-standing paradigms. This is, to our knowledge, the first foray into geothermal operations that has used reinforcement learning to make steamfield predictions and operational suggestions. Using the GOOML forecast model as a Reinforcement Learning (RL) environment, we have setup a recommendation engine that can be used to suggest changes to operational paradigms that could result in greater real-world efficiency of operations.

In these RL experiments, we allowed an autonomous agent to make a change within the baseline GOOML forecast model, observe the change, and receive a reward if the change resulted in increased power generation. A generalization of this process is illustrated in Figure 8. The agent was permitted to only enact changes to wells (representing wellhead pressure and resultant mass flow changes) and pipeline pressures.
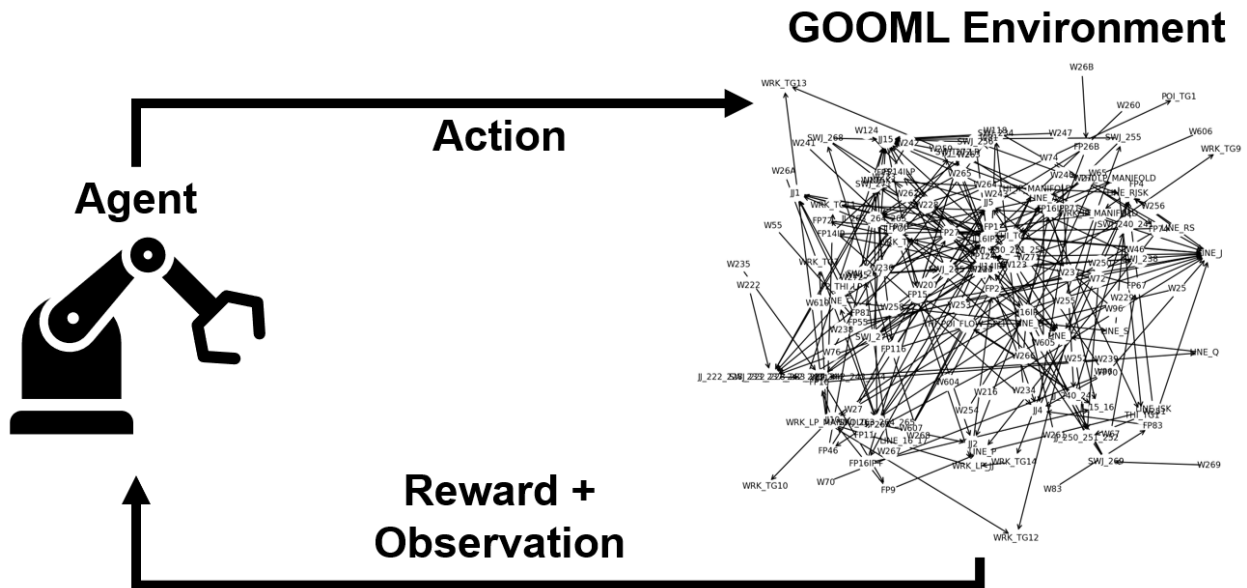


**Figure 8: Reinforcement Learning in the GOOML Environment. The agent can take an action in the GOOML environment- here represented by the spaghetti ball- (increase a well flow, throttle a control valve, etc.) and the result yields a reward or a penalty that affects total system power output.**

11

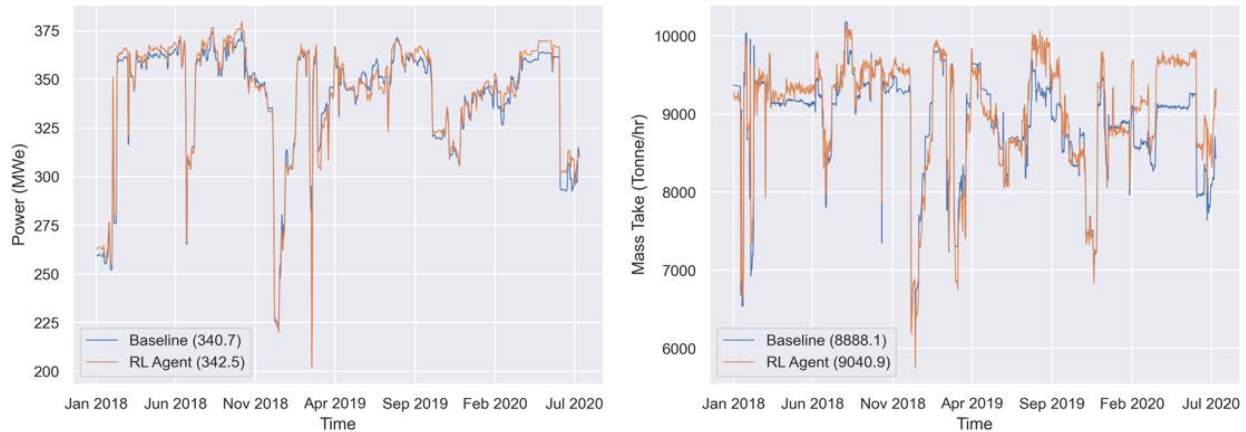| Experiment # | 1 | 2 | 3 |
|---|---|---|---|
| Agent Control Level | +/- 2 bar IP +/- 1 bar LP +/- 2 Bar WHP | +/- 2 bar IP +/- 1 bar LP +/- 2 Bar WHP | +/- 0.5 bar IP +/- 0.25 bar LP +/- 2 Bar WHP |
| Reward Center (Relative to Baseline) | 0 | 100 | 100 |
| Result | Failure! No Improvement over Baseline | 20 MWe average net gain | 1.8 MWe average net gain |



Figure 9: Results of Reinforcement Learning Experiments. Experiment 1 was not successful and was unable to achieve a greater overall output as the reward was set too low. Experiments 2 and 3 improved system output through a reward that was centered at 100% of the total baseline output. By aiming to achieve a higher than baseline, the agent was able to successfully achieve an overall higher production from the system.

In our first experiments, we did not set the rewards high enough and the agent was never able to exceed the baseline and the experiment was deemed a failure (Figure 9-1). Through setting a proper reward (higher output) the agent actively sought a better solution to the baseline and yielded our first successful experiment (Figure 9- 2A & 2B) where the agent was able to add up to 20 MWe on average over the period interrogated.

12

Subsequently, we sought a scenario for the agent controls (Figure 9 – 3A & 3B) that was representative of real-world conditions utilizing small pressure control changes and true to real-world field constraints. This scenario still resulted in an increase to the overall power generation for the field investigated whilst maintaining the comparative mass extraction constraints, as seen in Figure 10.



**Figure 10: A comparison of our baseline (forecast) model and the RL agent results showing an average increase in power by 1.8 MWe, this resulted in a bit higher mass extraction by 153 tonnes/hour.**

The RL experiments have many limitations and should be interpreted with caution. The RL agent has only an understanding of the steamfield as it exists in the GOOML system model. This includes all the assumptions in the system model, but also any modeling gaps that we have not yet captured. There are also risks of the RL agent taking the system model and component models to operating regimes for which we do not have training data. For example, by lowering the pressure of the steamfield system overall, the agent can extract additional steam mass, but in doing so it is potentially operating the flash plants and turbines at pressures that do not occur in the actual steamfield. One solution is to model the components using simplified theoretical models, such as modeling the flash plant separation process using only the thermodynamic steam quality at the separation pressure. This might allow for the model to be more accurate under a wider variety of out-of-sample test conditions.

More cross validation will need to be performed before the results of the RL experiments can be utilized directly in real-world operations. Despite the limitations of these explorations, the important takeaway is that the GOOML framework provides a testing ground with which operators and RL agents alike can explore the digital environment in ways that were not previously possible. The ability to run millions or even billions of steamfield analysis permutations is a powerful capability that should help improve operational efficiency in a large number of steamfields around the world.

13

## 6. Conclusions and Future Work

The GOOML digital environment allows a programmatic and machine-learning-driven approach to geothermal operational decision making. We have built a historical modeling framework that assimilates and ensures that clean datasets are available for training algorithms. We have also built a forecast modeling framework that can use trained models and simple seed data to predict future operations. By establishing these frameworks, we have built a series of tools that are useful for interrogating current and future geothermal operational decisions.

Our RL experiments demonstrate that the developed GOOML digital space is a powerful platform to investigate scenarios of operational conditions and system optimization. Future work now includes incorporating injection controls into our models, which will allow us to fully constrain the system both upstream and downstream. We intend to incorporate these constraints into the system which will allow a full system model to test operational scenarios. The power of this complete system will be found in both matching historical models and building system forecasts through a digital twin environment where scenarios can be rapidly changed and analyzed. Further, we intend to understand the power of reinforcement learning applications to geothermal operations. By understanding if our model can suggest true and realistic operating conditions, we may yet be able to unlock additional operational efficiency in real-world geothermal fields.

## REFERENCES

Church, E.F. "Steam Turbines" *McGraw-Hill Book Company, New York*. First Edition. (1928).

Siratovich, P., Blair, A., Weers, J. GOOML: Geothermal Operational Optimization with Machine Learning. GRC Transactions, Volume 44, 2020.