PROCEEDINGS, 51st Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, February 9-11, 2026
SGP-TR-230

Stochastic Thermo-Hydro Modeling and Neural Network Surrogate Development for Thermal Resource Assessment of the Galleries-to-Calories Geobattery

Victor Fakeye[1], Travis McLing[2], Christine Doughty[3], Yingqi Zhang[3], Patrick Dobson[3], Trevor Atkinson[2]

[1]Colorado School of Mines, 1500 Illinois St, Golden, CO 80401

[2]Idaho National Laboratory, 1955 Fremont Ave, Idaho Falls, ID 83415

[3]Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720

victorfakeye@mines.edu

Keywords: Thermal energy storage, District heating, Mines, Coupled process modeling, Machine learning

ABSTRACT

The Galleries-to-Calories Geobattery concept explores the use of abandoned coal mine workings for large-scale thermal energy transport and storage. The system involves injecting waste heat from a supercomputing facility into flooded mine galleries, where groundwater flow can store and transport thermal energy for potential recovery in downgradient district heating and cooling applications. To evaluate the feasibility and performance of the Geobattery under geological and operational uncertainty, we developed a suite of stochastic thermo-hydrological (TH) simulations using Monte Carlo sampling of key uncertain parameters (e.g., permeability, porosity, thermal conductivity, specific heat capacity) and operating conditions (e.g., injection rate, injection temperature).

Results identified injection rate and temperature as the most influential parameters governing thermal front propagation, while the geometry of the room-and-pillar structure played a critical role in directing the extent and orientation of thermal advancement. Optimal combinations of material properties for maximizing heat recovery were also determined. To address the high computational cost of coupled-process stochastic modeling, we trained a neural network surrogate model on 24,000 physics-based realizations, achieving an $R^2 > 0.99$ and MAE < 0.1 for temperature predictions at monitoring locations. This surrogate enabled an additional 100,000 realizations for global sensitivity analysis and probabilistic thermal resource assessment. The integrated stochastic physics–surrogate modeling framework offers a computationally efficient tool for quantifying uncertainty, identifying key drivers, and informing early-stage design decisions for Geobattery systems.

1. INTRODUCTION

As the global energy sector accelerates towards decarbonization, addressing the imbalance between energy demand and supply has become a critical challenge. The heating sector, which accounts for almost half of global final energy consumption, remains largely dependent on fossil fuels (IEA, 2022). Thermal energy storage technologies enable the long-term storage of heat to balance seasonal variations in supply and demand, while also supporting decarbonization using excess or waste heat sources in place of fossil fuels.

Abandoned flooded coal mines offer a unique opportunity to support district heating and cooling by leveraging their existing underground structures, extensive void space, and hydraulic connectivity to store and transport thermal energy. The Geobattery concept involves injecting heat, heat travels through the fluid, and is produced for heating and cooling (Figure 1). Several studies have been carried out on the usage of flooded coal mine for heat storage and transport (Watzlaf and Ackman, 2006; Loredo et al., 2016; Athresh et al., 2019; Alvarez et al., 2021; Perez Silva et al., 2022, Receveur et al., 2025).

The Galleries2Calories (G2C) project is an international consortium exploring the use of abandoned coal mines in Edinburgh, Scotland, to store and transport waste heat from the University of Edinburgh Advanced Computing Facility (ACF) for district heating and cooling. Previous studies on the G2C Geobattery concept have been conducted by research teams from the University of Edinburgh, Lawrence Berkeley National Laboratory and Idaho National Lab.

Todd (2023) developed a THM model to examine the effects of injecting hot water into the Midlothian mine working without considering groundwater flow. Doughty et al. (2024) developed a deterministic model to gain insight into the range of thermal performance of the Geobattery using a simplified representation of the through-going mine-water flow path. Atkinson et al. (2024) presented initial hydraulic simulation results, using tracers to map the convective thermal plume distribution without accounting for heat conduction. Zhang et al. (2025) investigated the potential thermal performance using a simplified conceptual representation of the system. While these studies provide valuable insights, most relied on simplified room-and-pillar geometries and did not incorporate geological complexities or structural features.

Building on the previous studies, this work presents a high-fidelity stochastic modeling approach that incorporates uncertainties in system and operational parameters, as well as the geometry of the room and pillar workings, enabling the assessment of a broad range of potential behaviors of the Geobattery operation. This approach supports improved performance prediction, operational design, and optimization of future deployment strategies.

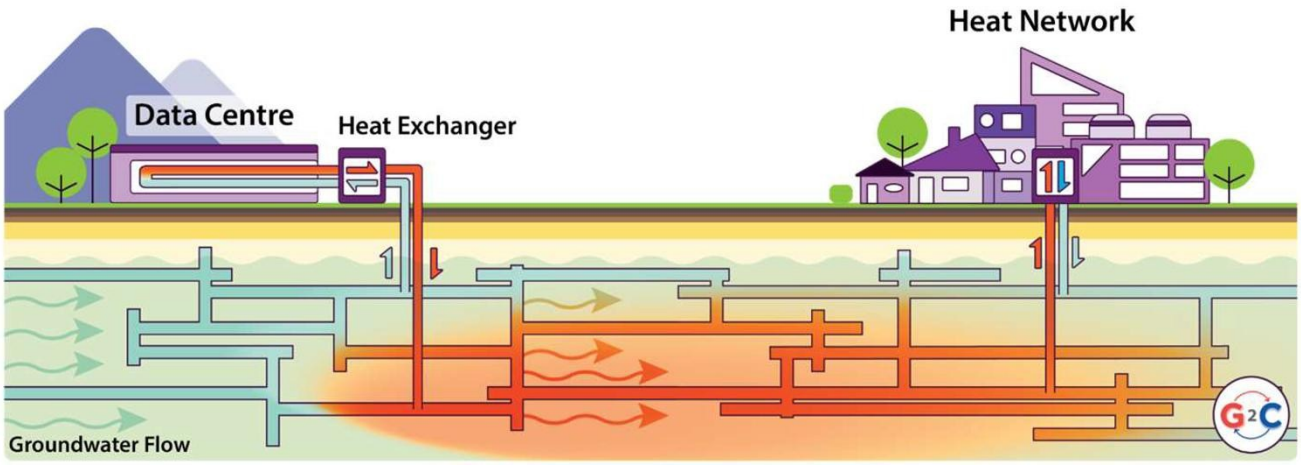

Figure 1: Geobattery concept illustrating data center waste heat injection into abandoned mines for downstream heat storage and delivery (Source: University of Edinburgh)

2. GEOLOGY AND HYDROGEOLOGY

The Geobattery site lies within the Midlothian Coalfield in the Midland Valley of Scotland (Figure 2), bounded to the southwest by Devonian–Silurian sedimentary and volcanic rocks making up the Pentland Hills. The regional structure is strongly influenced by the Pentland Hills Fault, which downthrows toward the southeast (Tulloch & Walton, 1958; British Geological Survey, 2003). The target coal seams, the Peacock and Kittlepurse, are inclined and laterally extensive and are separated by relatively low-permeability layers that can compartmentalize flow and focus groundwater movement along preferential pathways (Figure 3a) (Tulloch & Walton, 1958; British Geological Survey, 2003; Atkinson et al., 2024). Within the coalfield, the subsurface is dominated by Carboniferous successions of sandstones, siltstones, mudstones, limestones, ironstones, and coal-bearing strata, including the Lower Limestone, Limestone Coal, and Upper Limestone formations (Figure 3b). The regional hydrologic system is influenced by surface drainage through the North and South Esk Rivers, and groundwater recharge is supported by precipitation and runoff from surrounding uplands. Following mine closure, mine water rebound flooded the abandoned workings, forming a connected mine-water system in which flow can be guided by interconnected roadways, shafts, and hydraulically conductive flow paths. Consistent with this setting, groundwater movement is interpreted to be predominantly from Roslin toward Ramsay, supporting the concept of downgradient heat transport through the mine network (Todd, 2023; Atkinson et al., 2024). Reported mine-water temperatures in Scottish coalfields commonly fall in the range of ~12–21°C, with spatial variability influenced by depth, recharge source, and geothermal gradient (Gillespie et al., 2013; Farr et al., 2020, Receveur et al., 2025).

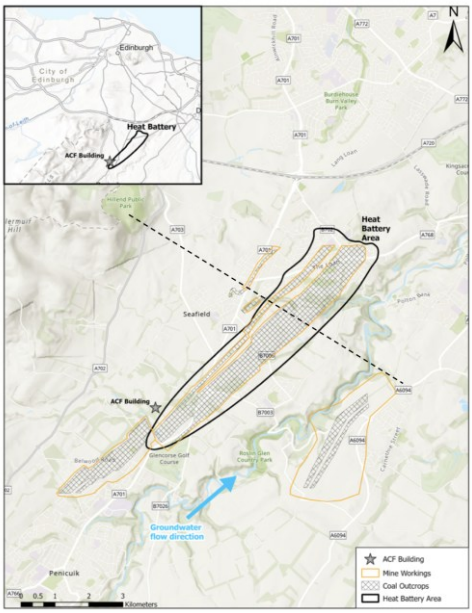

Figure 2: A map of the project site south of Edinburgh. University of Edinburgh's Advanced Computing Facility is marked as ACF on the map.

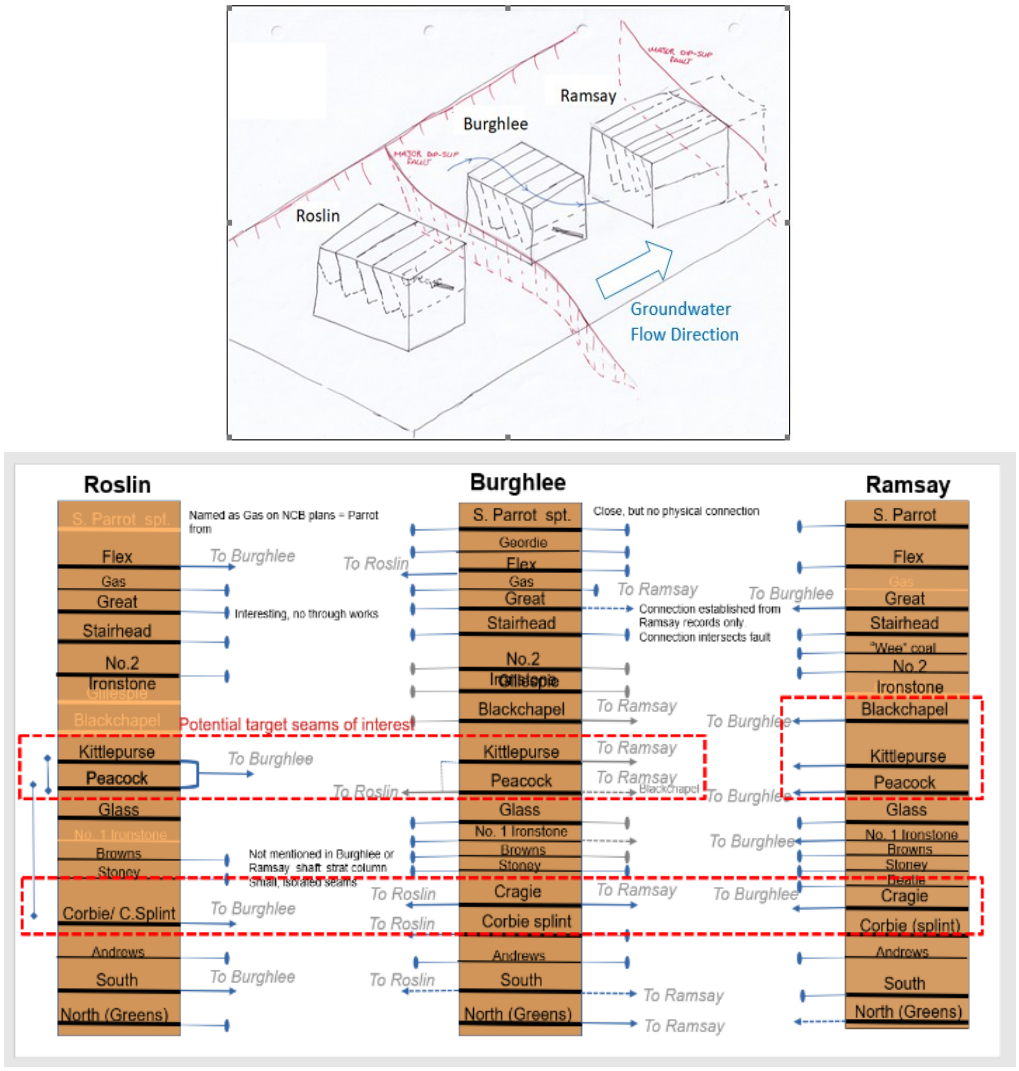

Figure 3: (a) The three collieries within the Geobattery study area and (b) target seams of interest within the three collieries, with red dashed lines indicating hydraulic connectivity (Atkinson et al., 2024; Doughty et al., 2024).

3. METHODOLOGY

This study integrates two complementary approaches: (i) high-fidelity physics-based numerical modeling for coupled thermo-hydro processes and (ii) machine-learning-based surrogate models to support predictive modeling of system performance. The overall methodological workflow is illustrated in Figure 4.

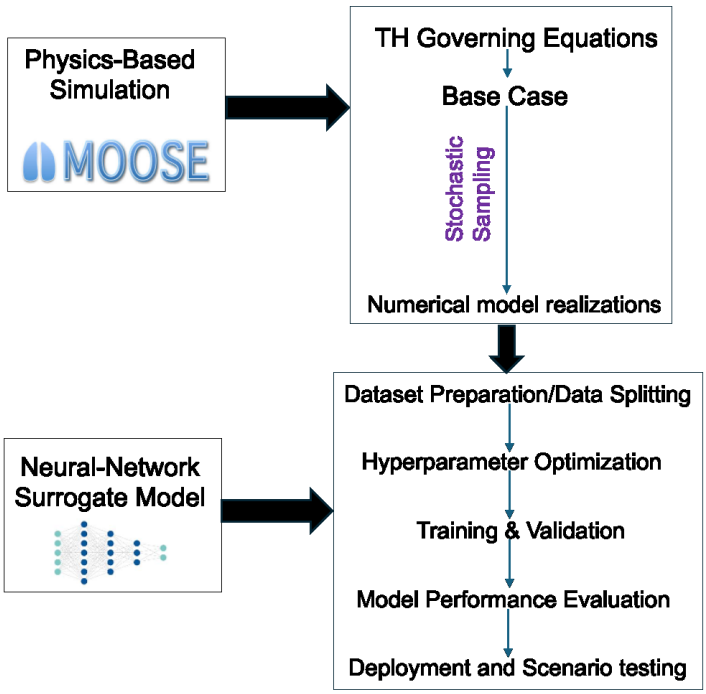

Figure 4: Workflow applied in this study

3.1 Numerical Modeling

The numerical simulations in this study were performed using the Multiphysics Object-Oriented Simulation Environment (MOOSE), an open-source finite element framework developed by the Idaho National Laboratory for solving coupled multiphysics problems (Gaston et al., 2009; Permann et al., 2020). We utilize the PorousFlow module designed for implementing fully coupled THM simulations in porous media (Wilkins et al., 2021).

3.1.1 Thermo-Hydro Governing Equations

The subsurface is represented as a saturated porous medium, and the governing equations describe the conservation of fluid mass and thermal energy.

Fluid Flow

$$\left(\frac{\dot{P}}{M} - A\dot{T}\right) - \nabla_i\left(k_{ij}\left(\nabla_j P - \rho g_j\right)/\mu\right) = 0$$

where P, T: fluid pore pressure, temperature, overdot represents a time derivative, M: Biot's modulus, A: effective volumetric thermal expansion coefficient, ρ: fluid density, gj: gravitational acceleration, kij: permeability tensor, μ: fluid viscosity, ∇: spatial derivative, ρR: density of rock grains, CR: specific heat capacity of the rock grains, Cv: specific heat capacity of the fluid, λij: thermal conductivity of the rock-fluid system.

This equation represents the conservation of mass for a slightly compressible fluid in a porous medium, where fluid density is temperature- and pressure-dependent. The first term accounts for fluid pressure changes due to compressibility and thermal expansion, while the second term represents Darcy flow driven by pressure gradients and gravity.

Heat Flow

$$\frac{\partial}{\partial t}\left((1-\phi)\rho_R C_R T + \phi\rho C_v T\right) - \nabla_i\lambda_{ij}\nabla T - \nabla_i\left(C_v T\rho\frac{k_{ij}}{\mu}\left(\nabla_j P - \rho g_j\right)\right) = 0$$

This equation governs energy conservation, combining conductive and advective heat transport. The first term represents changes in thermal energy, the second term describes heat conduction, and the third term accounts for heat advection by moving fluid.

3.1.2 Initial and Boundary Conditions

The model domain was initialized to reflect the in-situ thermal and hydraulic conditions of the Midlothian coalfield. Hydrostatic pressure was imposed at both the upstream and downstream boundaries to represent the natural static water column. A groundwater gradient of 0.04 was applied across the domain, and a groundwater flux of 10^{-6} m/s was applied at the upstream boundary to represent natural inflow conditions. The geothermal gradient was set to 25 °C/km, with a reference temperature of 16 °C at a depth of 168 m below ground level.

The two boundaries perpendicular to the groundwater flow direction, as well as the top and bottom boundaries, were specified as no-flow conditions for both heat and fluid, ensuring that transport occurred only along the primary flow path. Heat injection was represented by imposing a constant temperature of 32 °C at the designated inlet location corresponding to the University of Edinburgh Advanced Computing Facility.

3.1.3 Reservoir Model

Using available geological data for the Midlothian coalfield, the project team constructed a detailed geological model that incorporates key structural features, including faults, fractures, and stratigraphic layers (Figure 5). While this full-scale model captures the complexity of the site, meshing it in its entirety while preserving all fine-scale details would be computationally prohibitive. Therefore, a subset of the large-scale model was extracted for simulation purposes, retaining the essential geological and structural elements relevant to the Geobattery concept. The selected model encompasses the full spatial extent of the three collieries, which serve as the primary conduits for heat storage and transfer, as well as the major fault zone that separates two of the collieries. It also incorporates the mapped dip and thickness variations of the targeted coal seams (Figure 6a).

Given that the internal geometry of the historical mine workings is poorly constrained, the coal seams were subdivided into blocks representing the room-and-pillar configuration typical of such operations (Figure 6b). Material properties were assigned to these blocks using probabilistic distributions, enabling stochastic representation of uncertain parameters such as permeability, porosity, and thermal conductivity. This approach ensures that the model captures both the large-scale geological framework and the variability inherent in the subsurface system.

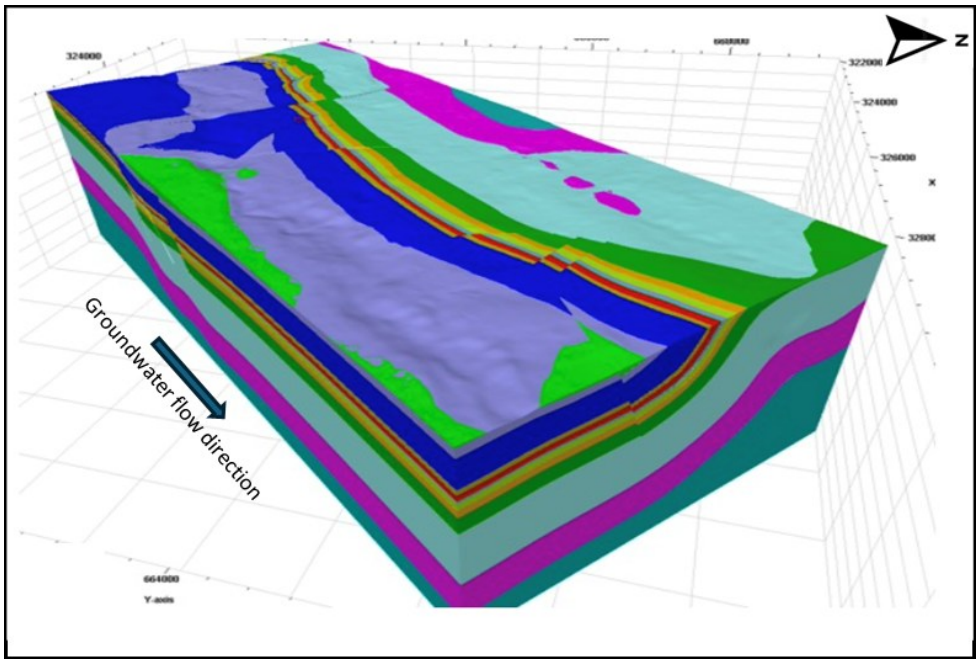

Figure 5: Geological model with detailed features (modified after Atkinson et al., 2024)

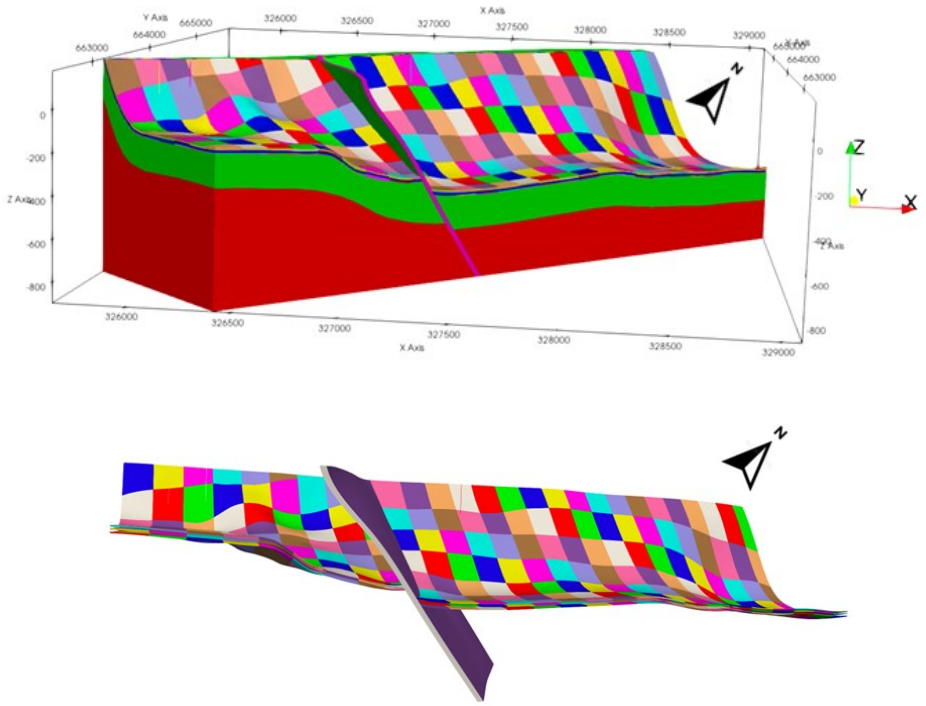

Figure 6: (a) Fully meshed model containing important features for numerical simulation (top) and (b) target coal seams subdivided into blocks with each color code representing rooms and pillars.

3.1.4 Base Case Parameters

Table 1 summarizes the model input parameters for the base case, derived from published data by Todd (2023), Atkinson et al. (2024), and Doughty et al. (2024). Permeability and porosity were assigned to reflect the hydraulic role of each domain. The mine rooms were treated as open void space (high k, $\phi \approx 1$) to represent efficient circulation and advective heat transport, while the coal pillars and host sandstone were assigned lower permeabilities and porosities to behave as relatively low-flow solids that store heat primarily through conduction. The fault zone was given an intermediate permeability to represent a potential preferential pathway relative to intact rock, whereas the bedrock was assigned very low permeability to act as a confining boundary and prevent unrealistic far-field leakage. Thermal properties were chosen to preserve realistic contrasts between water-filled rooms (high heat capacity, lower conductivity) and surrounding solids (lower heat capacity, higher conductivity), allowing the model to capture the dominant processes controlling thermal storage and propagation in mine workings.

Table 1: Model input parameters for base case.

Parameter	Host Rock (Sandstone)	Pillar (Coal)	Room	Fault	Bedrock
Permeability (m^2)	5e-13	1e-13	1e-9	1e-12	1e-18
Porosity (-)	0.215	0.055	1.0	0.07	0.01
Specific heat capacity (J/kg.K)	1050	1200	4184	1200	1050
Thermal conductivity (W/m.K)	3.25	1.35	0.6	1.35	3.25
Density (kg/m^3)	2500	2000	1000	2000	2500

3.1.5 Stochastic Modeling Framework

While the base case provides a deterministic reference, subsurface properties and operational parameters are inherently uncertain due to limited site characterization and natural heterogeneity. To account for this, we applied a stochastic modeling framework. Monte Carlo

simulations were performed to evaluate how uncertainty in input parameters propagates through the thermo-hydro model. 24000 realizations were run, each representing a unique combination of input variables sampled from their respective probabilistic distributions.

The primary objective of this approach was to generate a probabilistic range of thermal responses, enabling assessment of system performance variability and the spatial-temporal evolution of the thermal front. This methodology allowed for robust predictions that incorporate not just deterministic base case behavior, but also the plausible range of system states under realistic uncertainty. The uncertain parameters considered in the TH simulations, along with their probability distributions are presented in Table 2.

Table 2: TH parameters with uncertainty

Parameter	Distribution
Injection temp (C)	Normal (Mean = 35, SD = 5)
Injection rate (L/s)	Normal (Mean = 17, SD=3)
Thermal conductivity host rock (W/m.K)	Uniform (2 – 4)
Thermal conductivity pillar (W/m.K)	Uniform (0.2 – 1.5)
Specific heat capacity_host (J/kg.K)	Normal (Mean = 1000, SD=50)
Specific_heat_capacity_pillar (J/kg.K)	Normal (Mean = 1100, SD= 100)
Porosity_host (-)	Uniform (0.3 – 0.1)
Porosity_pillar (-)	Uniform (0.1 – 0.01)
Permeability_pillar (m^2)	Uniform (1E-17 – 2E-13)
Permeability_host (m^2)	Uniform (1E-16 – 1E-12)
Permeability_fault (m^2)	Uniform (1E-13 – 1E-11)

3.2 Surrogate Model Development

To address the high computational cost of running thousands of physics-based thermo-hydro simulations, this study develops a machine learning-based surrogate model using an Artificial Neural Network (ANN). The surrogate serves to significantly accelerate large-scale uncertainty quantification, enabling rapid predictions while preserving high accuracy.

ANNs are supervised machine learning algorithms capable of capturing complex nonlinear relationships between inputs and outputs. They consist of an input layer for receiving features, one or more hidden layers where nonlinear transformations are applied, and an output layer that generates the prediction. Each neuron within the network computes a weighted sum of its inputs, adds a bias term, and passes the result through an activation function. The network learns this weights and biases through backpropagation, an iterative optimization process that minimizes the prediction error as defined by a loss function (Figure 7). Because subsurface thermo-hydro processes are highly nonlinear, ANNs are well suited for modeling them.

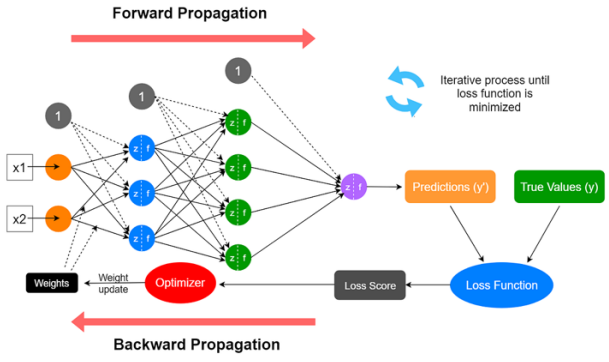

Figure 7: Learning process of a neural network (Pramoditha, 2022).

3.2.1 Dataset Preparation

The dataset used for the development of the ANN model consists of the 24000 realizations from the Monte Carlo simulations of the TH model based on the probabilistic distribution listed in Table 2. The input layers for the neural network represent the parameters with uncertainty (permeability, porosity, specific heat capacity, thermal conductivity) and the output layer represents the temperature increase at monitoring points. Exploratory data analysis was conducted to examine distributions and relationships between parameters. Following standard procedure, all inputs were normalized to a range between 0 and 1 using *Scikit-learn's* MinMaxScaler to ensure consistent scaling and comparable influence during model training. Splitting the dataset into training and testing subsets is essential for developing a robust machine learning model. The training set is used to train the model, while the testing set is used to assess how well the model generalizes to unseen data. For this study, we split the dataset into 80% training and 20% testing sets.

3.2.2 Hyperparameter Optimization

The predictive performance and generalization capability of an ANN are strongly influenced by its hyperparameters, which includes the number of hidden layers, number of neurons per layer, activation functions, learning rate, batch size, and regularization settings. Selecting suboptimal hyperparameters can lead to underfitting, where the model fails to capture the complexity of the underlying thermo-hydro (TH) processes, or overfitting, where the model captures noise and loses predictive power on unseen data.

To determine the optimal ANN configuration for this study, we employed a random search approach which involves sampling random combinations of hyperparameter combination from the specified distribution. The search space for the hyperparameter optimization is summarized in Table 3. The hyperparameter search was configured with a maximum number of trials set to 20, representing the total number of different hyperparameter combinations explored. Each configuration was trained three times to account for stochastic variations in weight initialization and data shuffling, ensuring a more robust evaluation of the hyperparameter configuration. The optimization objective used was the minimization of validation mean absolute error (val_mae). The validation set, comprising of 20% of the training data, was used to monitor the model's performance and guide the hyperparameter tuning process. Each candidate model was trained for up to 150 epochs with early stopping applied to monitor the validation loss. Training was halted if no improvement in validation loss was observed over 10 consecutive epochs. The optimal hyperparameter values obtained from the search consists of 3 hidden layers with properties presented in Table 4.

Table 3: Search space for hyperparameter optimization

Hyperparameter	Range
Number of hidden layers	1 – 5
Number of neurons per hidden layer	32 – 608 (step: 32)
Activation function	relu, tanh, sigmoid
Dropout rate	0.0 – 0.5 (step: 0.1)
Optimizer	adam, rmsprop, nadam, adamw

Table 4: Optimal hyperparameter values

Layer	Neurons	Activation	Dropout rate	Optimizer
Input	416	relu	-	-
Hidden Layer 1	192	sigmoid	0.1	nadam
Hidden Layer 2	32	relu	0.0	nadam
Hidden Layer 3	32	relu	0.0	nadam

3.2.3 Model Performance Evaluation

After identifying the best hyperparameters, the ANN was retrained using the combined training and validation datasets to maximize the amount of data available for learning. The finalized model was then applied to the unseen test dataset to generate temperature predictions at the monitoring locations. The predictive accuracy of the optimized ANN was evaluated using multiple performance metrics described below:

Mean Absolute Error (MAE) is the average of the absolute differences between the predicted values and the actual values. It measures the accuracy of continuous variables by calculating the average magnitude of the errors in a set of predictions, without considering their direction. Lower MAE indicates better model performance. It is easy to interpret since it uses the same units as the target variable.

Root Mean Squared Error (RMSE) is the square root of the average of the squared differences between the predicted values and the actual values. It penalizes larger errors more significantly than MAE due to the squaring of the error term. Lower RMSE indicates better model performance. It is more sensitive to outliers compared to MAE, due to the squaring of each error term.

Mean Absolute Percentage Error (MAPE) is a measure used to assess the accuracy of a regression model by calculating the average absolute percentage difference between the actual values and the predicted values. MAPE expresses the error as a percentage, which makes it easier to interpret the relative size of the error in comparison to the actual values. Lower MAPE indicates better model performance. It is scale-independent, making it useful for comparing performance across different datasets.

R^2, or the coefficient of determination, is a statistical measure that explains the proportion of the variance in the dependent variable that is predictable from the independent variables. It ranges from 0 to 1. $R^2 = 1$ indicates that the model perfectly explains all the variance of the target variable. $R^2 = 0$ indicates that the model does not explain any of the variance. Higher R^2 values indicate better model performance.

On the test data, the ANN achieved $R^2 > 0.99$, MAE < 0.1°C, and RMSE = 0.11, indicating excellent agreement with the physics-based TH model. Parity plots of ANN prediction versus physics-based simulation values showed data points closely aligned along the 1:1 line, and residual distributions were centered around zero, suggesting minimal systematic bias (Figure 8).

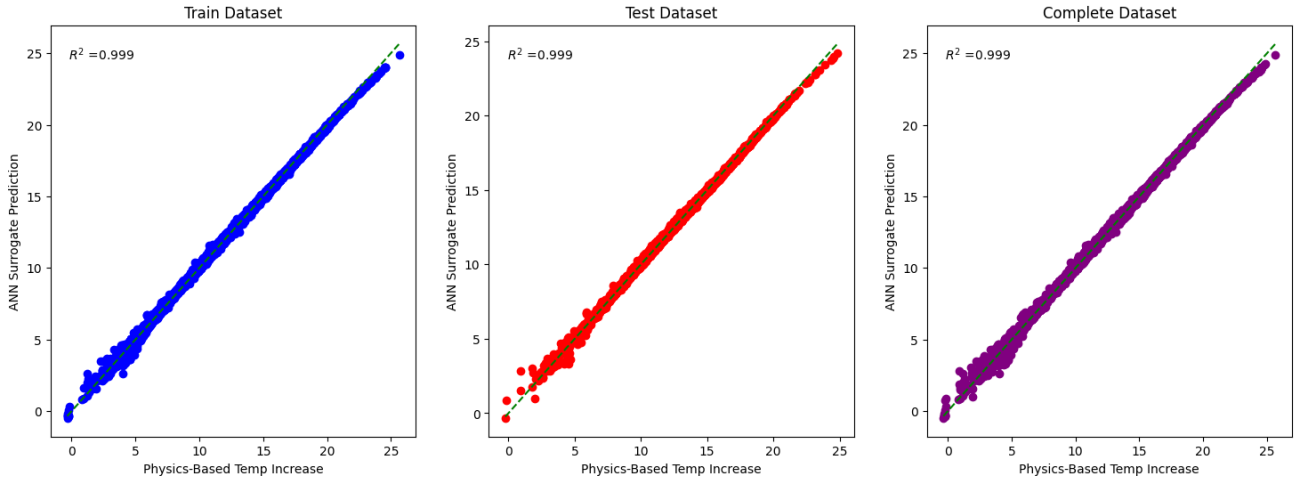

Figure 8: ANN surrogate model versus physics-based model.

4. RESULTS AND DISCUSSION

4.1 Base Case

Using the base-case parameters outlined in Table 1, Figure 9 presents the simulated temperature distribution within the coal seams after 10 years of continuous hot water injection. The results indicate that the thermal front has just propagated beyond the fault zone. Consequently, communities located upstream of the fault are likely to be the primary beneficiaries of the Geobattery operation during the initial decade of operation. L1 and L3 are monitor points to track the advancement of the thermal front.

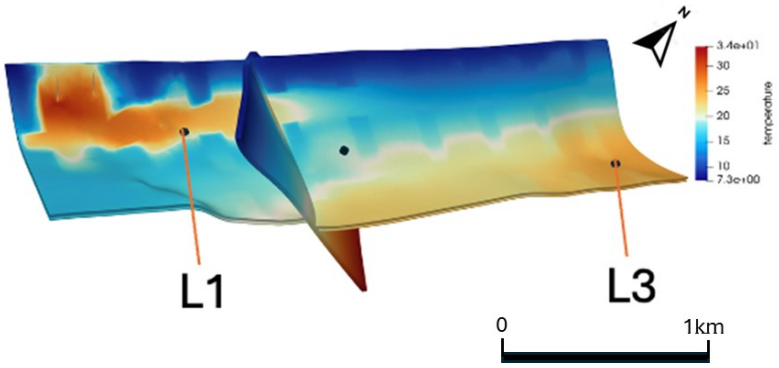

Figure 9: Thermal front propagation after 10 years of injection.

4.2 Effect of Pillar-Room Geometry on Thermal Front Propogation

The influence of pillar-room geometry was examined for 4 different configurations to evaluate both best- and worst-case scenarios (Figure 10). Across all cases, the primary control on thermal propagation is the degree of hydraulic communication between mine rooms, since heat transport is largely governed by advective circulation through connected pathways. The best-case scenario (configuration A) assumes a relatively intact and uniform room-and-pillar network with open cross-connections, allowing efficient flow redistribution and a broader, more uniform heated footprint. In contrast, increasingly restricted scenarios (configurations B & C) represent progressive loss of connectivity that produces channelized flow, localized heating along dominant pathways, and delayed warming in distal zones. These variations are physically plausible in legacy mine systems and may be caused by localized mine collapse/roof falls and debris accumulation, sediment or fines infill that partially seals corridors, or operational/equipment-related blockages such as restricted inflow/outflow sections, clogging, or limited access to certain entries. The worst-case configuration (D) reflects a highly restricted network where pillars effectively behave as barriers, which may occur under severe collapse, long-term closure, or persistent operational/equipment-related restrictions that strongly limit circulation and confine heat delivery near the primary flow path. Together, these scenarios provide realistic bounds on thermal performance under uncertain mine conditions and evolving connectivity.

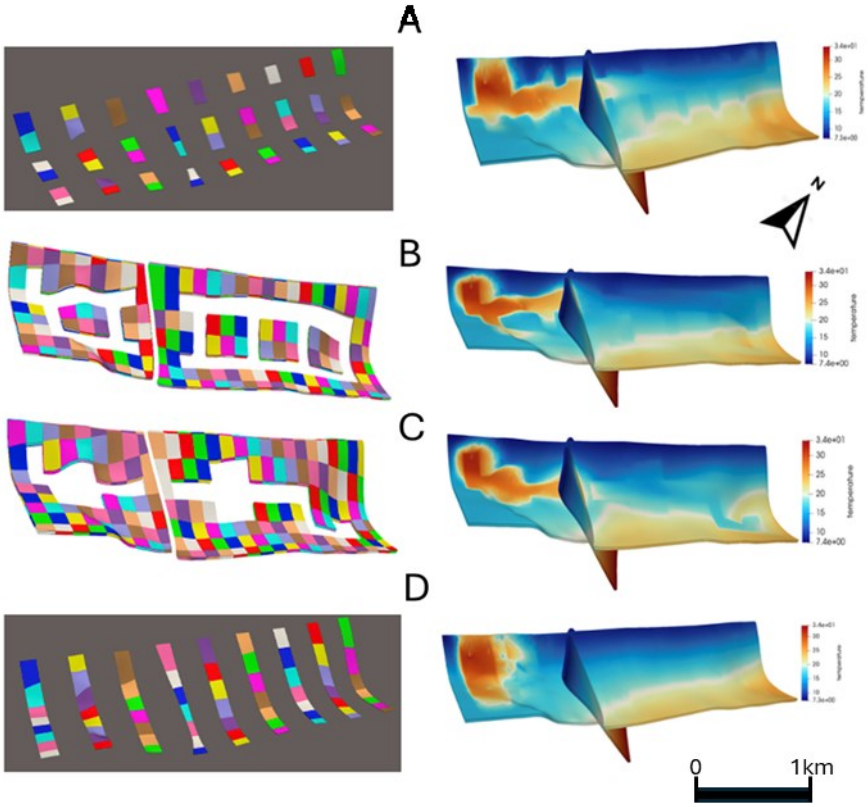

Figure 10: Thermal front propagation from different room and pillar configurations. Best case: top; worst case: bottom.

4.3 Uncertainty Quantification

The primary Quantity of Interest (QoI) for the 24000 realizations is the temperature increase at monitoring locations L1 and L3 after 10 years of continuous fluid injection, relative to their respective initial conditions. This metric quantifies the extent of thermal energy propagation to each location, thereby providing insights into the advancement of the thermal front under geological and operational uncertainty.

Figure 11 presents the resulting probability distributions for the temperature increase at L1 and L3. L1, located about 650m away from the injection point, exhibits a nearly symmetric, unimodal distribution centered around ~12–13 °C, with a standard deviation indicating moderate spread. This shape suggests that despite variability in input parameters, the thermal signal at L1 remains strong and relatively predictable. L3, positioned farther from the injection source, shows a broader, right-skewed distribution with a mean around 0.05 °C. The small temperature increase at L3 indicates that the thermal front is not anywhere close to L3 within 10 years.

These probability distributions provide a quantitative measure of prediction uncertainty at different monitoring points. They also indicate the degree of reliability in achieving thermal performance targets under the range of geological and operational conditions considered.

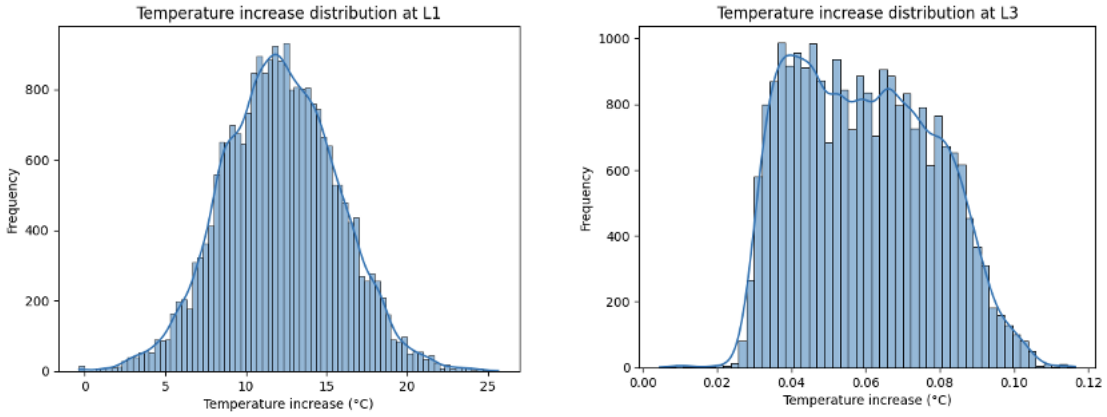

Figure 11: Probability distribution of temperature increase at L1 and L3.

4.4 Neural Network Validation

Following validation of the surrogate model, the ANN was deployed to predict outcomes for 100,000 parameter combinations sampled from the input parameter distribution presented in Table 2. This large-scale evaluation enabled comprehensive uncertainty quantification and global sensitivity analysis, allowing rapid exploration of parameter influence on thermal performance without the prohibitive computational cost of equivalent numerical simulations. As shown in Figure 12, the probability distributions obtained from the surrogate-model realizations reproduce the overall range, central tendency, and dominant trends observed in the 24,000 physics-based simulations at both monitoring locations. The agreement is particularly strong at L1, where the surrogate captures the full distribution shape and spread. At L3, the surrogate predictions remain consistent in magnitude and variability but exhibit a modest deviation in distribution shape relative to the physics-based results. This behavior is expected given that temperature increases at L3 are substantially smaller in amplitude and more sensitive to threshold- and regime-dependent transport behavior, such that small surrogate prediction errors can lead to visible changes in the distribution form. Overall, the close correspondence between surrogate and physics-based results supports the use of the ANN for efficient uncertainty propagation and sensitivity assessment, while also highlighting that far-field predictions may exhibit slightly increased distributional uncertainty.

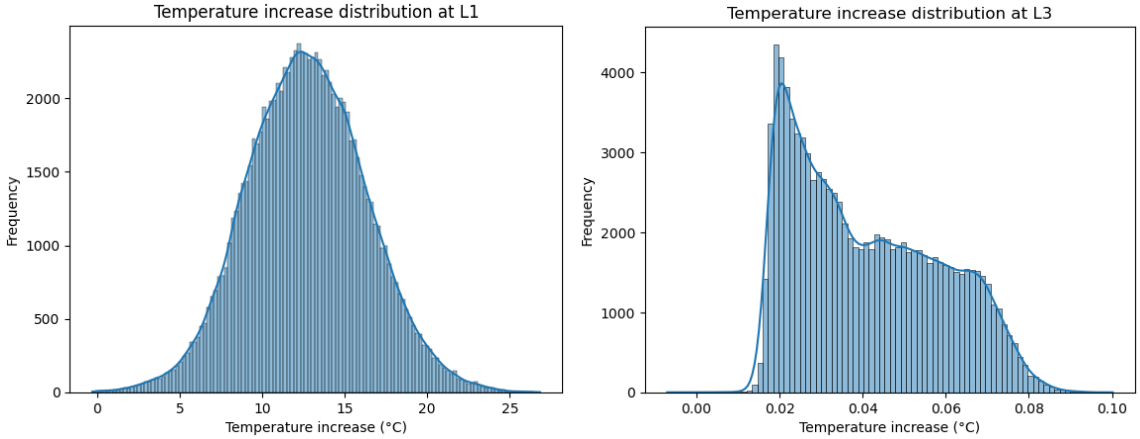

Figure 12: 100,000 realizations from trained surrogate model.

4.5 Sensitivity Analysis

A global sensitivity analysis was conducted to determine the relative contribution of each uncertain input parameter to the variability in temperature increase at L1 and L3. This helps to identify the key parameters influencing thermal front propagation and the optimal material properties for heat production, offering valuable insights into the feasibility of the Geobattery operations.

Two complementary approaches were used: Sobol variance-based sensitivity, computed directly from the physics-based TH model outputs shown in Figures 13 & 14 and ANN-based sensitivity, which was derived from the trained surrogate model, shown in Figures 15 & 16.

The first-order Sobol indices quantify the direct influence of each parameter on the output variance, while the total-order indices capture both the direct effects and interactions between parameters. For example, increasing the injection temperature directly increases the thermal energy carried by the injected fluid and therefore raises the predicted temperature change at the monitoring locations, while increasing the injection rate directly enhances advective heat transport and accelerates the propagation of the thermal front through the

system. Interaction effects occur when the influence of one parameter depends on the value of another, which is particularly relevant for coupled transport processes in heterogeneous pathways. For instance, the effect of injection rate on temperature rise at the downstream location (L3) can depend strongly on fault-zone permeability, since higher injection rates translate to meaningful downstream heat delivery only when the fault provides sufficient transmissivity for flow communication. Similarly, the impact of injection temperature may be amplified or suppressed by the pillar–room geometry, as better-connected configurations promote circulation and distribute thermal energy more efficiently, whereas poorly connected geometries can confine heat locally even under higher injection temperatures. These coupled behaviors explain why a parameter may exhibit a modest first-order index but a larger total-order index, reflecting its importance primarily through interactions that jointly control advective heat transport and thermal plume migration. Together, these methods provide a robust assessment of parameter influence, accounting for nonlinear dependencies and multi-parameter interactions inherent in the TH system.

The results from both ANN-based and Sobol sensitivity analysis indicate that:

-Injection rate and injection temperature are the most influential parameters, contributing the largest share of the variance in temperature rise at both monitoring locations.

-Pillar–room geometry significantly affects heat transfer efficiency, with higher connectivity between rooms enabling more uniform and rapid thermal front advancement.

-Permeability of the fault zones plays a secondary yet important role, influencing both the rate and extent of advective heat transport for L3, which is situated downstream of the fault. The permeability of the fault directly controls the volume of heated water reaching this point.

-At L3, the mean temperature increase is only 0.06 °C after 10 years, indicating that heat conduction is the dominant transport mechanism there, unlike the more advection-driven response at L1. Specific heat capacity and thermal conductivity have minimal impact in advection-dominated regions, though they become more relevant in conduction-dominated zones such as L3.

Overall, the sensitivity analysis confirms that operational controls (e.g., injection rate and temperature) should be prioritized in design optimization, while thermal property uncertainties are of lesser concern in predominantly advective zones but gain importance in conduction-dominated regions downstream of major flow barriers.

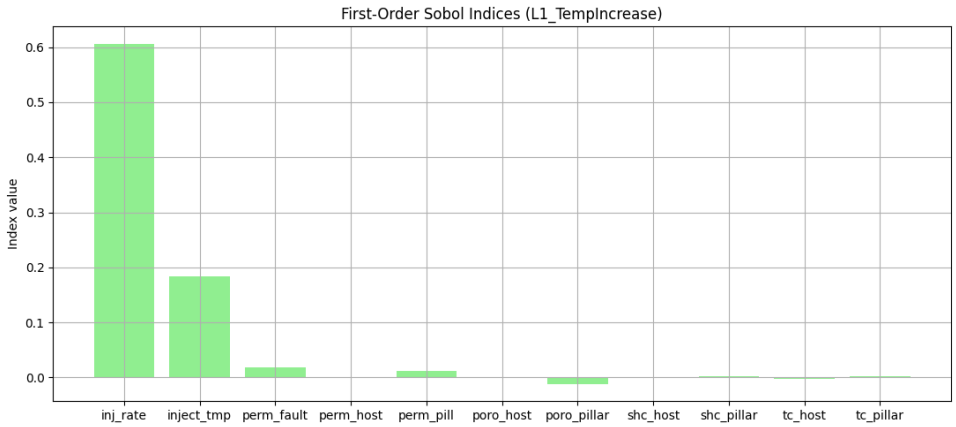

Figure 13: First Order Sobol Analysis at L1.

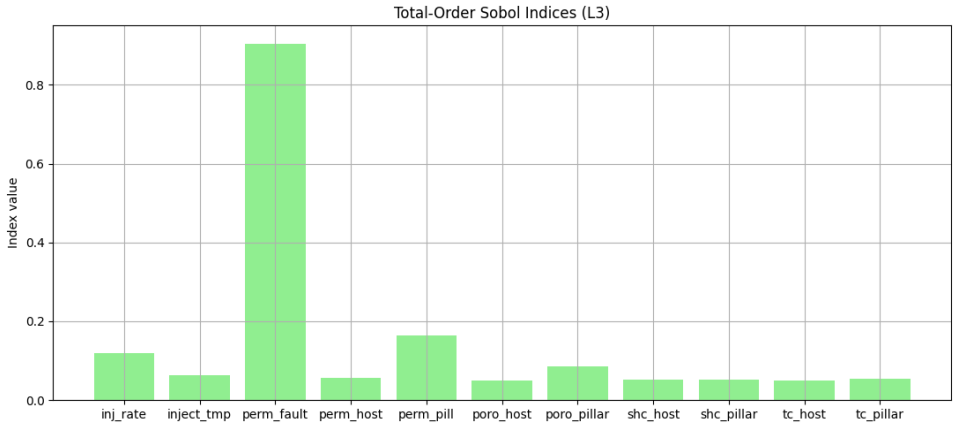

Figure 14: Total Order Sobol Analysis at L3.

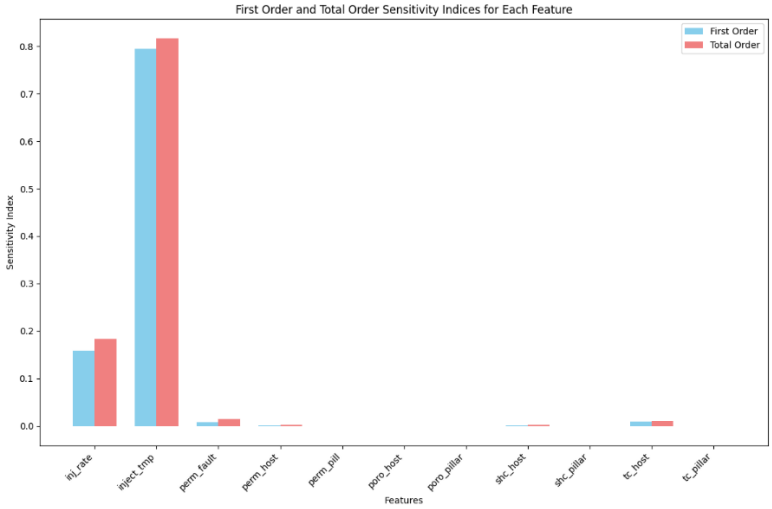

Figure 15: Surrogate-Model First and Total Order Analysis at L1.

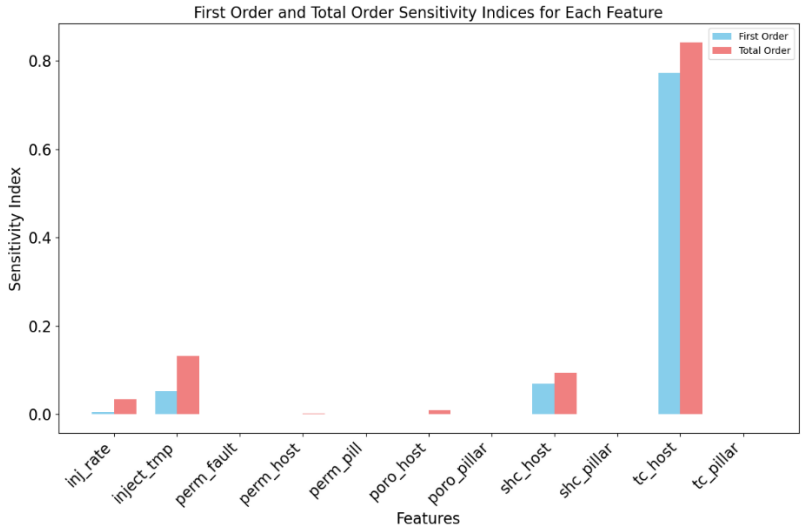

Figure 16: Surrogate-Model First and Total Order Analysis at L3

5. CONCLUSIONS AND FUTURE WORK

This study has explored a range of potential TH behaviors using a sophisticated coupled-process modeling approach, leveraging both physics-based and machine learning surrogates to evaluate 100,000 realizations. The analysis identified key parameters that are most critical for the thermal performance of the system.

For future work, field data will be collected to reduce uncertainty in the parameter space. One well has been drilled at the project site and data collection continues in order to characterize the subsurface and hydrogeologic flow regime. These data will be incorporated to constrain the stochastic input ranges, including boundary conditions, after which the reduced set of realizations will be used in coupled thermo-hydro-mechanical (THM) simulations. The THM modeling will be further refined by integrating surface settlement measurements to constrain mechanical responses. This sequential narrowing of uncertainty will yield a small, high-fidelity set of models that can be used to robustly assess the feasibility and operational performance of the G2C Geobattery.

ACKNOWLEDGMENTS

We acknowledge the G2C team including the University of Edinburgh, TownRock Energy Ltd., University of Strathclyde, University College Dublin, Scene Connect Ltd., and Sandown Ltd. for their collaboration and contributions that made this work possible. This research was supported by the U.S. Department of Energy (DOE), Office of Energy Efficiency and Renewable Energy (EERE), Geothermal Technologies Office, Low-Temperature Geothermal Program, under Award No. DE-AC02-05CH11231 (LBNL) and Contract No. DE-AC07-05ID14517 (INL). The views expressed are those of the authors and do not necessarily reflect those of the DOE or the U.S. Government. The G2C project is funded through the GEOTHERMICA consortium by Scottish Enterprise (Scotland), Geological Survey Ireland (Ireland), and the DOE (USA).

REFERENCES

Álvarez, H., Domínguez, G., Ordóñez, A., Menéndez, J., Álvarez, R., & Loredo, J. (2021). Mine water for the generation and storage of renewable energy: A hybrid hydro–wind system. International Journal of Environmental Research and Public Health, 18(13), 6758. https://doi.org/10.3390/ijerph18136758

Atkinson, T., McLing, T., Egert, R., Jin, W., Doughty, C., Dobson, P., Oldenburg, C., Zhang, Y., Kneafsey, T., McDermott, C., Graham, S., Perez Silva, A., & Rodriguez Salgado, P. (2024). Galleries-to-Calories (G2C): An International Collaboration Evaluating Thermal Energy Storage in Abandoned Mines for District Heating. GRC Transactions, 48, 1905–1914.

British Geological Survey. (2003). Edinburgh (Scotland) Sheet 32E: Bedrock geology, 1:50,000. British Geological Survey.

Doughty, C., Dobson, P., Oldenburg, C., Zhang, Y., Kneafsey, T., Egert, R., McLing, T., Atkinson, T., & Jin, W. (2024). Thermohydrogeologic modeling for the Geothermica project G2C (Galleries to Calories). In Proceedings of the 49th Workshop on Geothermal Reservoir Engineering. Stanford University.

European Commission: Directorate-General for Research and Innovation, Athresh, A., Al-Habaibeh, A., Banks, D., Zawartka, P., Burnside, N., Gzyl, G., Parker, K., Głodniok, M., Janson, E., Łabaj, P., Boyce, A., Younger, P., Roqueñí, N., Siodłak, Ł., Skalny, A., Loredo, J., García, A. G., Fernández, A. P., ... Hyria, A. (2019). Low-carbon after-life: Sustainable use of flooded coal mine voids as a thermal energy source: A baseline activity for minimising post-closure environmental risks (LoCAL): Final report. Publications Office. https://data.europa.eu/doi/10.2777/715020

Farr, G., Busby, J., Wyatt, L., Crooks, J., Schofield, D. I., & Holden, A. (2020). The temperature of Britain's coalfields. Quarterly Journal of Engineering Geology and Hydrogeology, 54(3), qjegh2020–109. https://doi.org/10.1144/qjegh2020-109

Gaston, D., Newman, C., Hansen, G., & Lebrun-Grandié, D. (2009). MOOSE: A parallel computational framework for coupled systems of nonlinear equations. Nuclear Engineering and Design, 239(10), 1768–1778. https://doi.org/10.1016/j.nucengdes.2009.05.021

Gillespie, M. R., Crane, E. J., & Barron, H. F. (2013). Study into the potential for deep geothermal energy in Scotland: Volume 2. British Geological Survey Commissioned Report CR/12/131. ISBN: 9781782569862.

International Energy Agency. (2022). World Energy Outlook 2022. https://www.iea.org/reports/world-energy-outlook-2022

Loredo, C., Roqueñí, N., & Ordóñez, A. (2016). Modelling flow and heat transfer in flooded mines for geothermal energy use: A review. International Journal of Coal Geology, 164, 115–122. https://doi.org/10.1016/j.coal.2016.04.013

Pramoditha, R. (2022, February 1). Overview of a Neural Network's Learning Process. Medium. Retrieved from https://medium.com/data-science-365/overview-of-a-neural-networks-learning-process-61690a502fa

Perez Silva, J., McDermott, C., & Fraser-Harris, A. (2022). The value of a hole in coal: Assessment of seasonal thermal energy storage and recovery in flooded coal mines. Earth Science Systems and Society, 2, 10044.

Permann, C. J., Gaston, D. R., Andrs, D., Carlsen, R. W., Kong, F., Lindsay, A. D., Miller, J. M., Peterson, J. W., Slaughter, A. E., Stogner, R. H., & Tonks, M. R. (2020). MOOSE: Enabling massively parallel multiphysics simulation. SoftwareX, 11, 100430. https://doi.org/10.1016/j.softx.2020.100430

Receveur, M., Gonzalez Quiros, A., Monaghan, A., Starcher, V., Walker-Verkuil, K., & Boon, D. (2025). Characterizing the thermal state and recovery potential of flooded coal mines using fibre optic Distributed Temperature Sensing methods. Papers presented to the

European Geothermal Congress 2025, Zurich, Switzerland, 7–9 October 2025 (Paper 92). European Geothermal Energy Council (EGEC).

Todd, F. (2023). Modelling the thermal, hydraulic and mechanical controlling processes on the stability of shallow mine water heat systems (PhD thesis). The University of Edinburgh.

Tulloch, W., & Walton, H. S. (1958). The geology of the Midlothian Coalfield. Memoirs of the Geological Survey of Scotland. Her Majesty Stationary Office.

Watzlaf, G. R., & Ackman, T. E. (2006). Underground mine water for heating and cooling using geothermal heat pump systems. Mine Water and the Environment, 25(1), 1–14.

Wilkins, A., Green, C. P., & Ennis-King, J. (2021). An open-source multiphysics simulation code for coupled problems in porous media. Computers & Geosciences, 154, 104820. https://doi.org/10.1016/j.cageo.2021.104820

Zhang, Y., Doughty, C., Dobson, P., Oldenburg, C., Kneafsey, T., McLing, T., Atkinson, T., & Jin, W. (2025). Performance assessment using numerical models for thermal energy storage and transport via underground mine workings. In Proceedings of the 50th Workshop on Geothermal Reservoir Engineering. Stanford University.