

Bayesian Optimization of MMW Beam Ablation Under Material Uncertainties

Adriano Casablanca^{1*}, Katerina Adamopoulou^{1*}, Franck Monmont², Nikos Nikiforakis¹

¹Laboratory for Scientific Computing, Cavendish Laboratory, Department of Physics, University of Cambridge, UK

²Quaise Energy, Cambridge, Massachusetts, US

*These authors contributed equally to this work

franck.monmont@quaise.energy

Keywords: MMW Ablation, Bayesian Optimization, Optimization under Uncertainty, Computational Physics

ABSTRACT

Millimeter Wave (MMW) ablation is a novel technology designed to reach drilling depths of 10-20 km by direct energy irradiation, with significant applications for the development of Enhanced Geothermal Systems (EGS). Achieving these depths is critical to recover heat from high enthalpy geothermal reservoirs, which are otherwise inaccessible with conventional drilling technologies. Further development of MMW ablation requires a detailed understanding of beam-rock interactions to identify optimal beam parameters. In this study, we apply Bayesian optimization to a computational simulation of the ablation process, aiming to derive best practices for effective drilling when subject to uncertainties in material thermophysical properties. The simulation solves the multiphase conservation of energy in the solid, melt and vapor phase while modeling the irradiation from the MMW beam source. In this work, uncertainties arise from the lack of knowledge of the thermophysical properties of the rock and from the spatial randomness of the polycrystalline structure of the rock. The stochastic problem is formulated as a mean-variance optimization problem, where we seek to find the optimum incident beam variables satisfying the expected rate of penetration whilst minimizing its variance. Solutions are identified using Bayesian optimization, where sampling locations are determined by an acquisition function that combines marginalized Expected Improvement and objective variance. In general, the methodology yields optimal beam parameters as a function of stochastic material properties, providing valuable information for the future deployment of MMW drilling technology.

1. INTRODUCTION

Geothermal heat is a carbon neutral energy source, capable of providing a stable base load power supply necessary to complement alternative renewable energy sources to ensure grid stability (Lund and Toth 2021, Hamm et al. 2019). However, current applications of geothermal energy are largely limited to hydrothermal heat sources and are as such limited by scale and geography. Accessing geothermal resources from deeper wells where the rock temperature can exceed 300° C would lead to a significant increase in energy density and potentially a magnitude increase in heat extraction compared to conventional geothermal sources (Cladouhos and Callahan 2023). These super hot rock (SHR) conditions occur everywhere on the planet, but require drilling deep into hard and abrasive basement rock formations where mechanical drilling suffers from lower rates of penetration and increased component wear (Brittenham et al. 1982). A compelling alternative is to use high-powered Millimeter Waves (MMW) generated by a gyrotron to melt and vaporize rock, with distinct advantages in drilling speed at increased depths and pressure (Oglesby et al. 2014, Houde et al. 2020). So far, MMW drilling has been developed in laboratory conditions and significant work has been directed to bring the technology to field trials: this includes borehole stability analysis (Porlles et al. 2024), transmission line design (Wang et al. 2020), material property measurements (Shteinman et al. 2024), and fluid flow analysis for SHR geothermal reservoirs (Sani et al. 2012).

Computational models have been developed to model the drilling process as they offer a simpler and complementary alternative to expensive and time-consuming experimental tests. A numerical model was first developed by Oglesby et al. (2020) involving a single-phase heat transfer model. This work was significantly improved by Zhang et al. (2023), who developed a thermal model capable of capturing additional temperature-dependent behaviors such as phase changes and material removal. These responses are dependent on estimates for material behavior, which are obtained from data that can lack density or accuracy. In addition, our understanding of these properties decreases with depth due to increased pressure and temperature, hence, quantifying the impact of uncertainty on drilling performance is essential. In this study, the physical properties of the material and the distribution of the polycrystalline structure of the rock are treated as environmental variables, which capture conditions outside the users control. The effects of uncertainty in the drilling process can be divided into systematic errors in thermophysical properties and aleatoric uncertainty associated with the spatial polycrystalline structure of the rock. Variations in the thermophysical properties are accounted for by randomly sampling these quantities from a user-specified probability distribution while the grain spatial arrangement is obtained by randomly generating Voronoi tessellations which satisfy a given grain volume distribution.

In this study, we are interested in finding the best beam parameters that correspond to a target Rate of Penetration (RoP), while taking into consideration the impact of the environmental variables on the simulation output. As these uncertainties propagate to the output, the simulation results follow a probability distribution whose mean and variance can be computed (Girard et al. 2003). The properties of the resulting distribution are used in a mean-variance optimization framework, with the aim of finding a solution satisfying the expected target RoP with the smallest variance possible. Indeed, solutions exhibiting lower variance indicate a lower sensitivity to uncertain material

properties and are more robust, presenting less operational uncertainty and potentially decreasing total drilling time by minimizing downtime. Conversely without penalizing the variance, the optimization could converge to a minima with higher variance (Makarova et al. 2021). Incorporating variance as a measure of risk in the objective has been explored in Cakmak et al. (2020), Hong et al. (2014), Makarova et al. (2021), Sani et al. (2012), and Iwazaki et al. (2021).

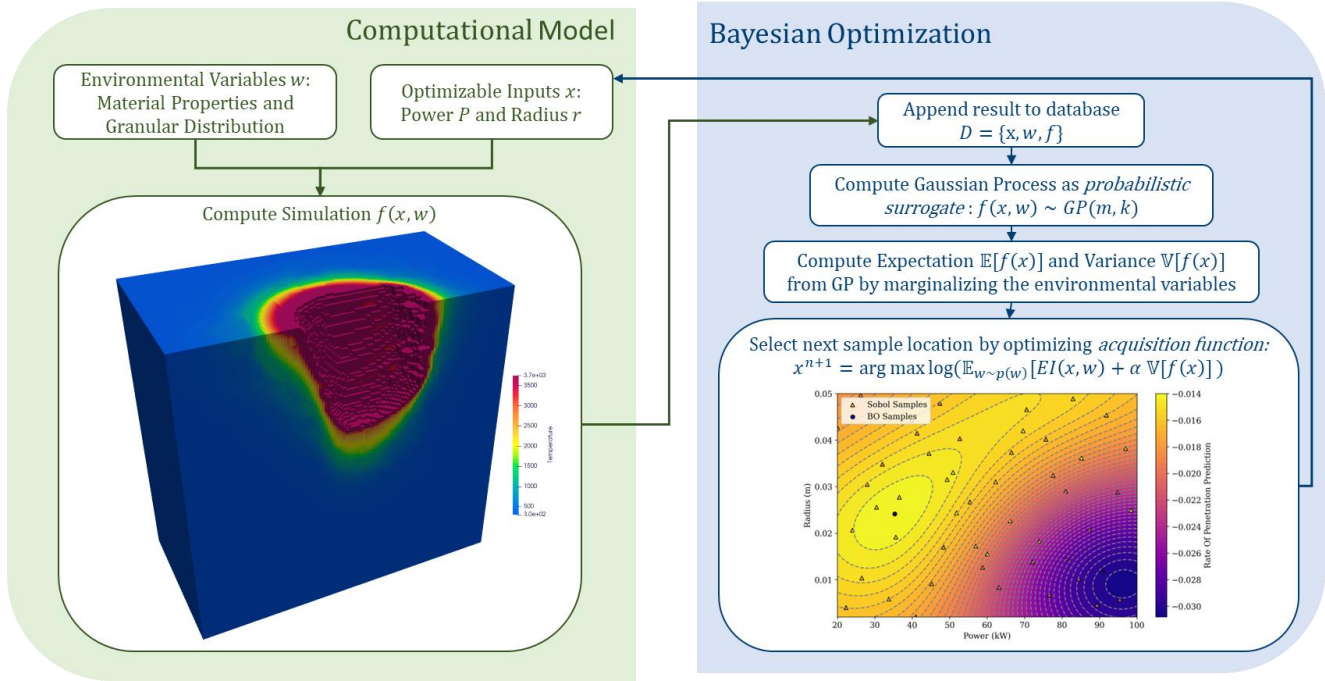


Figure 1: Mean-variance optimization of the rate of penetration is achieved by coupling a computational simulation of the MMW beam with Bayesian Optimization. The computational model takes as inputs the optimizable parameters, power and radius, as well as environmental variables which are randomly selected at each evaluation, and computes the predicted rate of penetration by evolving an enthalpy-based partial differential equation. Bayesian optimization takes the results of the simulation, and by first creating a probabilistic surrogate and then computing an acquisition function, guides the optimization loop by selecting inputs for subsequent evaluations of the computational model which perform well on average while providing low variance.

A Bayesian Optimization (BO) algorithm is used in this work as presented in Garnett (2023). It is an extrinsic framework that uses a probabilistic metamodel of the computationally expensive simulations to build a surrogate of the objective function, and an acquisition function to guide the optimization with a suitable policy to select the next set of design inputs. This modularity of BO provides great flexibility to optimize unique and complex objectives as the individual components can be tailored specifically to the problem at hand. First, the surrogate can be adapted to the function being modeled by employing a wide-variety of Bayesian models (Li et al. 2023); as a starting point, this work employs Gaussian Processes (GPs) (Williams and Rasmussen 2006). The GP model is trained on the beam variables and most influential environmental parameters. The environmental parameters are then marginalize out to find the expectation and variance of the objective function, as in El Amri et al. (2023). The acquisition function is also adapted to the specific optimization problem, where for mean-variance optimization the next sample location is chosen by balancing the marginalized Expected Improvement against the marginalized variance.

The following study is organized as follows: the optimization problem is described in Section 2, followed by the computational model and the environmental variables in Section 3. In Section 4, the Bayesian optimization under uncertainty algorithm is explained. The results are then presented in Section 5. The propagation of variable uncertainties to the model output, the rate of penetration, are examined first. Our Bayesian optimization algorithm is then applied to the optimization of the rate of penetration where power is the only input variable, and the latent heat of vaporization is the only source of uncertainty. Next, a high-dimensional optimization is presented where power and beam waist are the input variables, and 48 environmental parameters (influencing density, conductivity, specific heat, phase transition temperatures, absorptivity, reflectivity, and emissivity) are chosen as the sources of uncertainty.

2. OPTIMIZATION OBJECTIVES

We consider the problem of finding the best beam parameters minimizing the expected difference between the Rate of Penetration (RoP) and a target RoP. For a Gaussian beam profile, the RoP is defined as the maximum hole depth drilled at a target borehole radius divided by the total simulation time, as illustrated in Figure 2(b). To start with, beam power P and beam waist w_0 are chosen as input variables x . These variables are sufficient to characterize a beam with a Gaussian profile. A more realistic multi-modal beam would require many more design variables including the stand-off between the end of the transmission line and the rock. In addition, some environmental input parameters, w , used in the physical model are assumed to be random variables, because they are only known to a certain degree of

accuracy. For example, measuring the evaporation temperature and the corresponding latent heat of evaporation is extremely difficult and significant variability in the published data exists.




Figure 2: (a) Example simulation result for a heterogeneous material sample. The left half shows the percentage of vaporized volume fraction and the right shows the temperature of the material. (b) Example simulation results for a homogeneous material sample. Given a target radius for the wellbore, the wellbore depth is calculated as the maximum depth with a cross-section of the target radius. The rate of penetration is then computed by dividing the wellbore-depth by elapsed drilling time.

The resulting simulation output $f(\mathbf{x}, \mathbf{w})$, is consequently stochastic with a range of outcomes for identical beam settings. However, in the optimization study, only the influence of the beam parameters is of interest, and consequently the influence of the environmental variables on the model response need to be marginalized out. This is achieved by marginalizing the model response $f_{\mathbf{w}}$ with respect to the environmental parameters \mathbf{w} as follows:

$$f_{\mathbf{w}} = \int f(\mathbf{x}, \mathbf{w})p(\mathbf{w})d\mathbf{w}. \quad (1)$$

Finally, the optimization problem can be set as a Mean-Variance (MV) problem:

$$\arg \min_{\mathbf{x} \in X} \{ \mathbb{E} |f_{\mathbf{w}} - f^{target}| \} - \alpha \mathbb{V} \{ |f_{\mathbf{w}} - f^{target}| \}, \quad (2)$$

where X is the input variable space and α is a user-defined coefficient of risk tolerance. The first term of equation 2 is the objective to minimize, i.e. the expected difference between the target RoP f^{target} and the predicted RoP $\mathbb{E}|f_{\mathbf{w}} - f^{target}|$. The second term is the variance to be minimized. Solutions which perform well in a mean variance optimization, balancing performance and risk (Makarova et al. 2021), should lead to improvements in the drilling process.

3. MMW ABLATION SIMULATION

A computational model of the ablative process is used to model the RoP of the drilling process for given specific beam parameters. Drilling through rock using a MMW beam has three types of material removal: vaporization, melt flow, and thermal spallation. Vaporization is the most appropriate mode of removal to capture in this study as it can be modeled with an energy conservation equation at a low computational cost, while the melt flow requires complex variable density hydrodynamical models which are object of future study. In addition, thermal spallation is largely negligible as most of the drilling occurs above liquidus temperature. Modeling the vaporization process requires describing the first-order effects of the beam and simulating the absorption of the energy beam, the change in material properties, and the eventual vaporization of the geological material. The simulation code is written in AMReX, an open source C++ based software framework supporting Adaptive Mesh Refinement (AMR), which allows for dynamically increasing grid resolution in areas of interest (Zhang et al. 2019).

3.1 Enthalpy Partial Differential Equation

Vaporization of the geological materials is captured using a computer simulation built on the work by Zhang et al. (2023) coupled with an enthalpy-based partial differential equation. Alexiades et al. (2010) propose tracking energy changes through the volumetric enthalpy H , where infinitesimal changes are defined as:

$$dH = \rho c_p dT + (1 - \alpha_t) dp, \quad (3)$$

where ρ is the density, c_p the specific heat at constant pressure, T the temperature, α_t the thermal coefficient of expansion, and p the pressure. For computational efficiency a constant pressure assumption is made such that $dp = 0$. Hence, the energy equation can be described by treating volumetric enthalpy as a conserved variable:

$$\frac{\partial H}{\partial t} - \nabla \cdot \left(\frac{\kappa(T)}{\rho(T)c_p(T)} \nabla T \right) = q(T), \quad (4)$$

with t as time, κ the thermal conductivity, and the source term $q(T)$ as the power input Q from the MMW and the heat losses from radiation and convection:

$$q(T) = Q - (\varepsilon \sigma (T^4 - T_0^4) + h(T - T_0)) \delta, \quad (5)$$

where ε is the surface emissivity, σ is the Stefan-Boltzmann constant, and h the air convection coefficient, T_0 as the surface temperature, and δ as a Dirac delta function to compute the heat losses only at the melt front surface. The incident power of the beam is computed as:

$$\dot{Q} = I(r) \alpha(T) (1 - R) \exp(-\alpha(T)(z_0 - z)), \quad (6)$$

where α is the absorption coefficient, R the material reflection index, and $z_0 - z$ is the distance of the waveguide from the surface. I_0 is the beam intensity at the melt surface, dependent on the radial distance from the beam center r . The beam profile examined in this study is the Gaussian beam, with intensity profile:

$$I(r) = \frac{2 P_0}{\pi w_0^2} \left(\frac{w_0}{w(z)} \right)^2 \exp\left(\frac{-2r^2}{w(z)} \right), \quad (7)$$

where P_0 is the input power, w_0 as the beam waist radius, and $w(z)$ as the spot size radius:

$$w(z) = w_0 \sqrt{1 + \left(\frac{\lambda(z_0 - z)}{\pi w_0^2} \right)^2}. \quad (8)$$

3.2 Material Properties

Applications of the MMW beam are aimed towards use in hard basement rock formations, of which a principle component is granite. Granite is a granular igneous rock composed mainly of quartz, feldspar, and plagioclase; the exact composition varies by granite type, but typically contains more than 20% quartz and a ratio of 1.0 to 6.5 of plagioclase to feldspar (Bonin et al. 2020). Material properties of these components are modeled by temperature dependent polynomials. These are given with temperature T in degrees Celsius:

$$\rho(T) = a_\rho T + b_\rho, \quad (9)$$

$$\kappa(T) = a_\kappa + b_\kappa T + c_\kappa T^2, \quad (10)$$

$$c_p(T) = a_{c_p} T + b_{c_p}, \quad (11)$$


$$\alpha(T) = a_\alpha + b_\alpha \exp\left(\frac{-c_p}{T+273.15} \right). \quad (12)$$

Determining the exact value of these coefficients is difficult due to the high temperatures which need to be reached, and as such can only be identified with large ranges of uncertainty. The properties in Table 1 represent a current best-estimate for their value for quartz, feldspar, and plagioclase, obtained by regressing experimental data onto polynomials. However, the available data is often limited by range and sparsity, resulting in the coefficients carrying some corresponding uncertainty. This variability is captured in this model by drawing from the material property from a normal distribution centered around the above values with a standard deviation as some percentage of its value. For example, if the material property variance is determined to be equivalent to 10% of its value, then the b_ρ coefficient for feldspar is given as $b_\rho \sim \mathcal{N}(2700, 270)$.

Table 1: Material properties for Quartz, Feldspar, and Plagioclase.

<u>Material Dependent Properties</u>				<u>Material Independent Properties</u>	
Coefficient	Quartz	Feldspar	Plagioclase	Coefficient	All Materials
$a_p [kg/K]$	-0.0075	-0.0075	-0.0075	$a_\alpha [m^{-1}]$	252.71
$b_p [kg]$	2650	2700	2850	$b_\alpha [m^{-1}]$	62455.03
$a_\kappa [W/m K^3]$	17.5	2.5	2.5	$c_\alpha [K]$	9998.36
$b_\kappa [W/m K^2]$	-0.035	-2.75e-3	-8.5e-4	$\Delta L_{vap} [J/kg]$	13664000
$c_\kappa [W/m K]$	2.0e-5	1.85e-6	4.0e-7	$T_{vap} [K]$	3354
$a_{c_p} [J/kg K^2]$	0.3	0.4	0.4	$R [-]$	0.2
$b_{c_p} [J/kg K]$	850	650	750	$\varepsilon [-]$	0.8
$T_{melt} [K]$	1700	1200	1300		
$\Delta L_{melt} [J/kg]$	150000	362500	250000		

3.3 Granular Tessellation

**Figure 3: Example Voronoi tessellation visualized in 3D.**

The polycrystalline structure of granite is modeled using a 3D Voronoi tessellation approach (Bourne et al. 2022, Kuhn et al. 2020). The Voronoi diagram is constructed by defining a set of seed points $\{x_i, i = 1, \dots, n\}$, with each seed generating a cell that encompasses all points closer to it than to any other seed. Each cell is then randomly given a material and its corresponding properties according to the ratio of quartz, feldspars, and plagioclase of the chosen granite specimen.

The volume and shape distributions of the grains are approximated by packing spheres in 3D space, whose radii, number and allowed overlap are selected to approximately match the observed volume distributions of granite grains. These parameters affect the tessellation formed and tuning them allows the representation of a range of granite microstructures. This granite modeling approach enables us to investigate how mineral composition and granular configuration of granite influence drilling speed, offering insights into the interaction between rock properties and beam performance

4. METHODOLOGY: BAYESIAN OPTIMIZATION UNDER UNCERTAINTY

Bayesian Optimization (BO) is a type of optimizer specialized for global, complex, and expensive-to-evaluate functions. By taking a Bayesian approach, BO builds a surrogate to predict the objective function while providing uncertainty estimates, which it then employs in the acquisition function to guide the optimization process by balancing evaluations in regions of high uncertainty (exploration) against regions with high estimated function value (exploitation) (Garnett 2023). Since computing the surrogate and acquisition function can be

computationally expensive, BO is best suited to cases where the cost of evaluating the objective function remains the dominating factor. By maximizing the informational value of previous evaluations, BO achieves high sample-efficiency, where less samples are required to achieve high evaluation scores, leading to an overall lower cost of optimization. BO is managed through BoTorch, a Python-based library supporting a large variety of applications for BO (Balandat 2020). In the following paragraphs we describe the BO algorithm and highlight how uncertainty is handled in the optimization method.

4.1 Surrogate: Gaussian Processes

Gaussian Processes (GPs) are a Bayesian machine-learning model well-suited for BO due to their interpretability, computational efficiency for medium sized data sets, and well-calibrated uncertainty bounds (Garnett 2023, Williams and Rasmussen 2006). GPs are a collection of infinite random variables of the model response, such that the joint distribution of every finite subset is a multivariate Gaussian distribution of the k -dimensional random input vector \mathbf{x} . They are denoted by $f(\mathbf{x}) \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where the mean and covariance functions are defined as:

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad (13)$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]. \quad (14)$$

Selecting the covariance function $k(\mathbf{x}, \mathbf{x}')$ (also known as the kernel) is a critical decision, since it determines its ability to adjust to the training data. The kernel selected for this study is the so-called Matérn 5/2 kernel, which is widely used for its stationary properties and limited differentiability, making it well-suited for modeling physical processes:

$$k_{\text{Matérn-5/2}}(\mathbf{x}, \mathbf{x}') = \sigma^2 \left(1 + \frac{\sqrt{5}|\mathbf{x}-\mathbf{x}'|}{\ell} + \frac{5|\mathbf{x}-\mathbf{x}'|^2}{3\ell^2} \right) \exp\left(-\frac{\sqrt{5}|\mathbf{x}-\mathbf{x}'|}{\ell}\right), \quad (15)$$

where σ^2 is the variance parameter controlling the vertical scale of the kernel, and ℓ is the length-scale parameter governing the smoothness and horizontal scale of the function. Considering a training set $D = \{\mathbf{x}_i, \mathbf{f}_i\}_{i=1}^N$, the joint distribution of the training outputs \mathbf{f}_t and prediction outputs \mathbf{f}_* is given by:

$$\begin{bmatrix} \mathbf{f}_t \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K(\mathbf{X}_t, \mathbf{X}_t) & K(\mathbf{X}_t, \mathbf{x}_*) \\ K(\mathbf{x}_*, \mathbf{X}_t) & K(\mathbf{x}_*, \mathbf{x}_*) \end{bmatrix}\right). \quad (16)$$

The posterior predictive distribution is then obtained by applying the standard rule for conditioning which preserves normality:

$$\mathbf{f}_* | \mathbf{X}, \mathbf{f}, \mathbf{X}_* \sim \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\sigma}_*^2), \quad (17)$$

where the mean and variance are expressed as:

$$\boldsymbol{\mu}_* = K(\mathbf{X}_*, \mathbf{X}_t)[K(\mathbf{X}_t, \mathbf{X}_t)]^{-1}\mathbf{f}_t, \quad (18)$$

$$\boldsymbol{\sigma}_*^2 = K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X}_t)[K(\mathbf{X}_t, \mathbf{X}_t)]^{-1}K(\mathbf{X}_t, \mathbf{X}_*). \quad (19)$$

4.1.1 Accounting for Uncertainty in Environmental Variables

The results of the MMW simulation runs are kept in a dataset $D = \{\mathbf{x}, \mathbf{w}, \mathbf{f}\}$, from which a GP is constructed such that $f(\mathbf{x}, \mathbf{w}) \sim GP$. In practice, we are only interested in $\mathbf{f}_w(\mathbf{x})$, where we aim to marginalize the environmental variables as set out in the integral Equation 1. $\mathbf{f}_w(\mathbf{x})$ does not necessarily follow a Gaussian Process, but its distribution can be modeled by an expectation $\mathbb{E}[\mathbf{f}_w(\mathbf{x})]$ and variance $\mathbb{V}[\mathbf{f}_w(\mathbf{x})]$ calculated through the total law of iterated expectation and conditional variance (Girard et al. 2003). Using Monte Carlo (MC) sampling, the law of total expectation is:

$$\mathbb{E}[\mathbf{f}_w(\mathbf{x})] = \mathbb{E}_{w \sim p(w)}\{\mathbb{E}[f(\mathbf{x}, \mathbf{w})]\} = \mathbb{E}[\boldsymbol{\mu}] \approx \frac{1}{N} \sum_{i=1}^N \boldsymbol{\mu}_i^{MC}, \quad (20)$$

and the law of total variance:

$$\mathbb{V}[\mathbf{f}_w(\mathbf{x})] = \mathbb{E}_{w \sim p(w)}\{\mathbb{V}[f(\mathbf{x}, \mathbf{w})]\} + \mathbb{V}_{w \sim p(w)}\{\mathbb{E}[f(\mathbf{x}, \mathbf{w})]\} = \mathbb{E}[\boldsymbol{\sigma}^2] + \mathbb{V}[\boldsymbol{\mu}] \approx \frac{1}{N} \sum_{i=1}^N \boldsymbol{\sigma}_i^{2MC} + \frac{1}{N-1} \sum_{i=1}^N (\boldsymbol{\mu}_i^{MC} - \mathbb{E}[\boldsymbol{\mu}])^2, \quad (21)$$

where $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ are the GP mean and variance of the i th Monte Carlo sampling point.

4.2 Acquisition Function: Mean-Variance Optimization

The goal of BO is to find the global optimum in a domain whilst minimizing the number of expensive simulation calls. Given an initial data set, the next input vector \mathbf{x} is sampled directly from the acquisition function. The acquisition function uses the predictions and uncertainty estimates of the Gaussian Process to calculate the utility of sampling in a certain location by trading-off exploration (sampling in unexplored regions) and exploitation (sampling in promising regions). A commonly used acquisition function is the so-called Expected Improvement (Mockus 1975) which computes the point maximizing the expected improvement over the current optimum \mathbf{x}^* :

$$EI(\mathbf{x}) = \int_{-\infty}^{\infty} I(\mathbf{x}) \phi(z) dz = \int_{-\infty}^{\infty} \max(f(\mathbf{x}) - f(\mathbf{x}^*), 0) I(\mathbf{x}) \phi(z) dz, \quad (22)$$

where ϕ is the normal cumulative density function. In this work we seek the vector input \mathbf{x} minimizing the mean-variance problem, under uncertainty on the environmental variable \mathbf{w} . Hence, the acquisition function needs to be marginalized with respect to the specified distribution of \mathbf{w} . Marginalization is calculated via Monte Carlo integration, where a realization of \mathbf{w} is sampled repeatedly from $\mathcal{N}(\boldsymbol{\mu}_w, \boldsymbol{\sigma}_w)$. Therefore, our Marginalized Expected Improvement (MEI) acquisition function is:

$$\mathbb{E}_{w \sim p(w)}[EI(\mathbf{x}, \mathbf{w})] \approx \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} EI(\mathbf{x}, \mathbf{w}_j). \quad (23)$$

To guide the optimization process towards samples which perform well under a mean-variance objective, we define the Mean-Variance marginalized Expected Improvement acquisition function:

$$\mathbf{x}_* = \arg \max_{\mathbf{x} \in X} \log(\mathbb{E}_{w \sim p(w)}[EI(\mathbf{x}, \mathbf{w})] - \alpha \mathbb{V}[f_w(\mathbf{x})]), \quad (24)$$

which combines the MEI acquisition function with a variance term to penalize solutions with high variance. In addition, we take the logarithm of the function as it results in an identical optimization policy, but provides an expression which is easier to optimize numerically (Ament et al. 2023). While the Expected Improvement focuses on regions that are likely to yield better objective values (exploitation), the variance term, given by the sum of two terms of equation 21, favors both sampling in regions where the predictions of the model are more uncertain (exploration), and sampling in regions where the variations of the environmental parameters cause relatively small output fluctuations. This dual objective ensures that the optimization process not only seeks to improve the current best solution but also gathers valuable information in less-explored regions with low risk, thereby increasing the robustness and global convergence capability of the Bayesian Optimization algorithm.

5. RESULTS

In this section, we first explore the effects of uncertainties associated with the environmental parameters on the simulation outputs and the consequential impact on the mean-variance optimization problem. Next, a two-dimensional test case is presented, where the input variable is power and the latent heat of vaporization is the environmental parameter. Finally, an optimization is performed where the input variables are the beam power and waist, while treating the full set of material properties as environmental variables.

5.1 Distribution of outcomes due to material uncertainty

Introducing uncertainty by varying the physical properties and distribution of the rock granular structure produces simulation outputs that are no longer deterministic. In this section, the impact of environmental variables is analyzed by running multiple simulations with the same beam design variables, whilst varying environmental parameters. Simulation results are first obtained for three different test cases by varying the mean grain size while keeping the material property constant. In the second test, the mean grain size is kept constant and the uncertainty associated with the physical properties is varied. In both cases, the simulation is repeated 25 times and the distribution of grains is recomputed using a Voronoi tessellation. A normal distribution is fitted to the outputs, as in Figure 4.

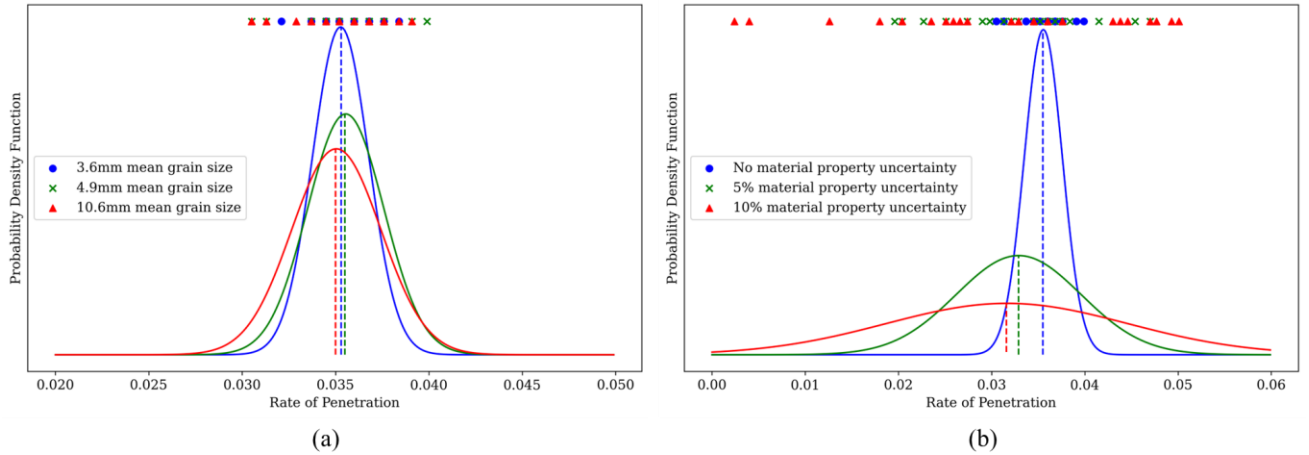


Figure 4: Probability density function of the model response. For each case, 25 model responses are computed (plotted at the top of the figure) and the normal distribution is fitted to the data. (a) Varying mean grain size. (b) varying material property uncertainty.

The first test where the grain size distribution is varied is shown in Figure 4(a). These results demonstrate that the variance of the model response increases with increasing mean grain size. The second test where the physical properties of the material are varied is illustrated in Figure 4(b). These plots show that greater uncertainty on the physical properties results in greater variance in the model response. Furthermore, when the uncertainty on the physical property increases, the mean model response, RoP, decreases. When the mean grain size of the rock increases, the variance of the model response increases too, due to the coarseness of the rock structure.

5.2 Bayesian optimization where beam power is used as input design variable and uncertainty is simulated by varying the latent heat of vaporization

We apply the BO framework to an optimization test where the only design variable, \mathbf{x} , is the beam power and the only environmental variable, \mathbf{w} , is the latent heat of vaporization. Two cases are considered, one with lower uncertainty on the knowledge of the environmental parameter and one with higher uncertainty. The material specimen modeled is a polycrystalline granite, where the volume fraction of quartz, feldspar and plagioclase volume fractions are fixed at 0.4, 0.3 and 0.3 respectively. The simulation is composed of 2000 Voronoi points, such that the average grain radius is 4.9 mm. A domain size of $[10 \times 10 \times 10] \text{ cm}^3$ with a coarse grid of 64^3 cells and 1 layer of AMR is implemented which doubles the resolution where the temperature is greater than 1600°C

The objective of the optimization is to find the beam power which minimizes the difference between a target RoP of 0.07 m/min with a target radius of 0.01m after 60 seconds and the predicted mean RoP, with a risk-tolerance coefficient α of 1. This is achieved by first running an initial set of 40 quasi-random space-filling Sobol samples for a power range of 20–100 kW and variations in the latent heat of vaporization of 6000–10000 kJ/kg. The remaining material parameters are taken from Table 1 and the beam-waist radius is set at 0.02 m. The next 15 samples are selected by sampling the Mean-Variance marginalized Expected Improvement acquisition function as explained in Equation 24. The input parameter \mathbf{x} is determined by the optimizer, and the latent heat of vaporization is sampled randomly from a Gaussian distribution with mean $\mu = 6500$ and variance σ^2 . Two cases are examined with $\sigma^2 = 5 \cdot 10^4$ and $\sigma^2 = 1.5 \cdot 10^5$ respectively.

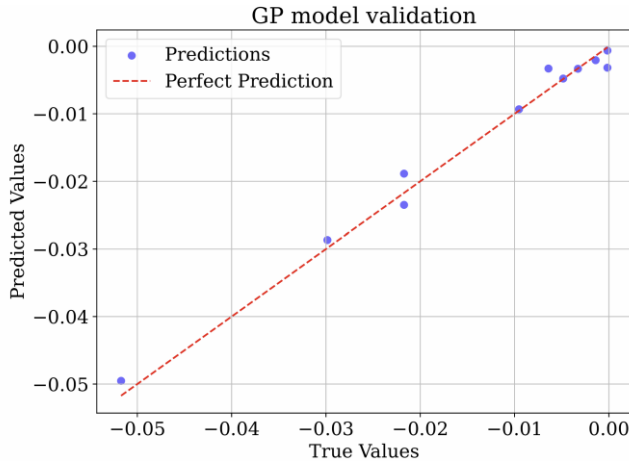


Figure 5: Prediction of the GP against the true RoP of the simulation for 10 randomly chosen test samples. The closer the predictions to the red dashed line, the better the agreement.

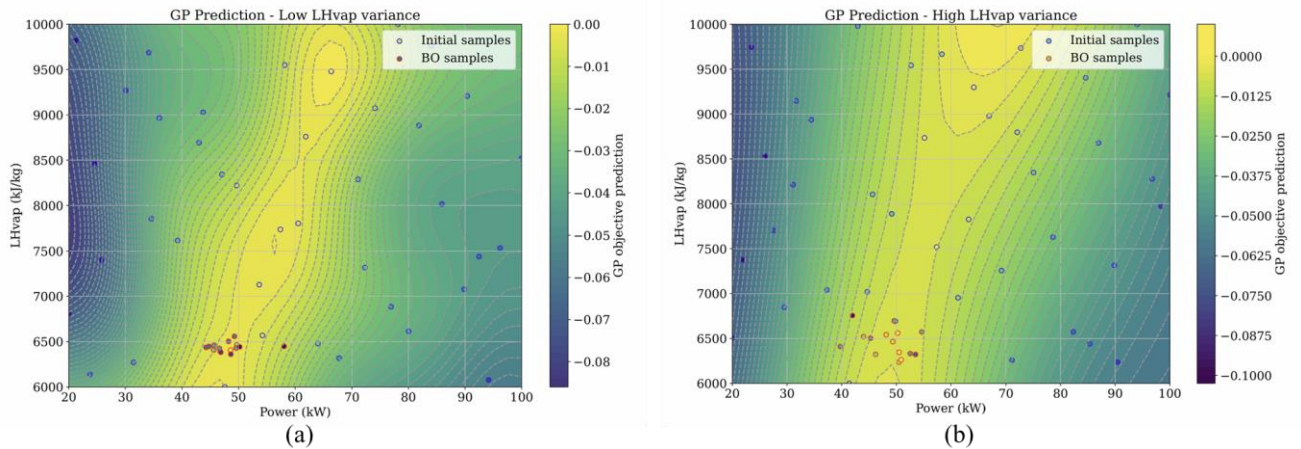


Figure 6: Results of BO run for two different values for the variance of the latent heat of vaporization. The contour plots display the predicted values of the objective function. (a) $\sigma^2 = 5 \cdot 10^4$ (b) $\sigma^2 = 1.5 \cdot 10^5$.

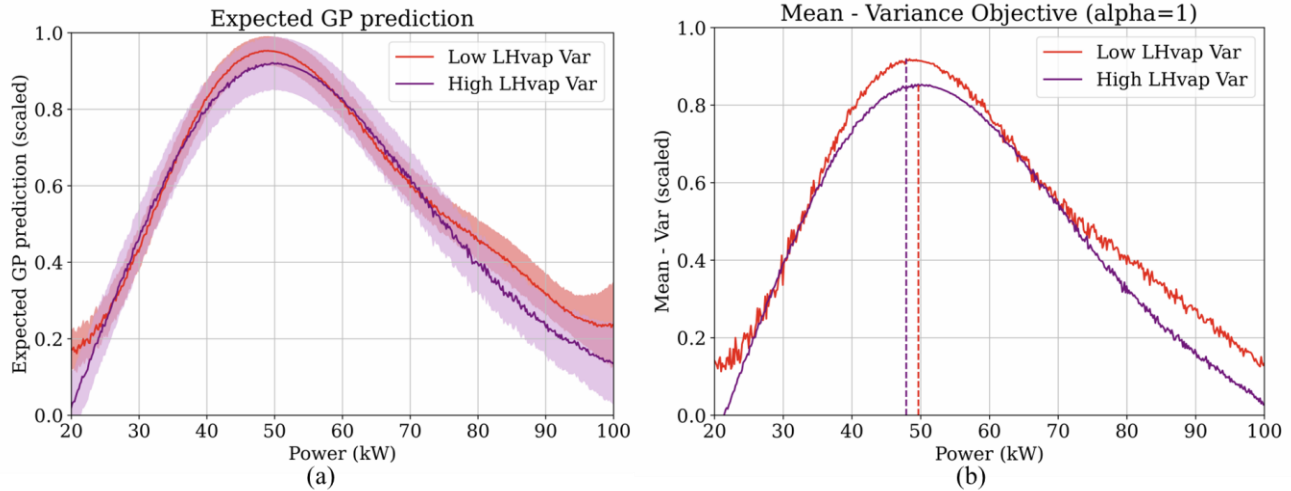


Figure 7: Marginalized GP prediction from the samples collected during the optimization study. The two test cases are the low latent heat of vaporization (LHvap) variance test with $\sigma^2 = 5 \cdot 10^4$ and the high variance case with $\sigma^2 = 1.5 \cdot 10^5$. The noise in the curves are a result of the Monte Carlo marginalization approach which is computed from a limited number of samples. (a) Mean GP prediction against power. The predicted variance is represented by the width of the shaded region. (b) Mean-Variance against power for $\alpha = 1$. The dashed line represents the maxima for the curve.

The accuracy of the GP predictions is first illustrated in Figure 5 where the surrogate predictions are compared against the simulation outputs for 10 samples. The results of the BO study are presented for the two test cases in Figure 6. As expected, the GP model predicts that a higher power is required to achieve the same RoP when the latent heat of vaporization increases. The BO samples for both tests tend to converge to the region minimizing the mean RoP prediction. However as shown in Figure 6(b), when there is a higher variance in environmental variable w , this can lead to a higher variance in the selection of the tunable parameter x .

The expected objective and the mean-variance function are plotted against power in Figure 8, using the GP model built with the complete sample-set in both cases. In 7(a), the expected objective and its expected variance are illustrated. They are computed using Monte Carlo marginalization on the w variables. They are then combined to calculate the Mean-Variance function as plotted in Figure 7(b). The results indicate that a larger variance in w leads to a higher expected variance, with a larger maximum than the lower variance case. In both scenarios, the optimal beam power is approximately 50 kW, specifically 48.0 kW for the low variance case and 49.8 kW for the high variance case. Hence, a higher variance in the latent heat of vaporization results in an increase in the required power to achieve the mean target RoP. This is most likely due to a compensating effect against the high latent heat of evaporation cases which result in much lower RoP. This interpretation is further supported by the asymmetry of the curve of the expected objective in Figure 7(b). It is steeper on the left hand side of its peak compared to its right hand side, indicating a larger sensitivity to beam power as the power decreases.

5.3 Bayesian optimization where beam power and beam waist are used as input design variables and uncertainty is simulated by varying all the physical parameters of the numerical model

Next we consider the optimization problem with power and radius as input design variables \mathbf{x} whilst varying the complete set of environment variables in Table 1. The min and max values of beam power and beam waist vary between 20 to 100 kW and 0.002 to 0.05 m respectively. A 10% standard deviation was assigned to the material parameters. The target RoP is fixed at 0.02 m/min with a target radius of 0.01 m. The numerical simulations are performed with 1 layer of AMR on top of a baseline coarse grid resolution 64^3 for a domain size of $[10 \times 10 \times 10] \text{ cm}^3$. The volume fractions of quartz, feldspar and plagioclase are equal to 30%, 60%, and 10% respectively. The parameter space is first explored with 50 quasi-random Sobol samples, followed by a BO loop of 10 evaluations. The optimization process is performed three times: first two times with a BO algorithm using the marginalized GP and Mean-Variance marginalized Expect Improvement with a risk tolerance coefficient of $\alpha = 1$ and $\alpha = 10$ as shown in Figure 8, and thirdly with a BO algorithm using a GP not trained on the environmental variables, using the Expected Improvement acquisition function as shown in Figure 9. Due to the high-dimensionality of the environmental variables (48 parameters), we marginalize only a limited number of environmental variables since GPs can struggle to learn relationships in a high-dimensional space. Only the most influential material properties are kept in the marginalization process. These properties are: conductivity coefficient a_k , the latent heat of melting ΔL_{melt} , and the latent heat of vaporization ΔL_{vap} .

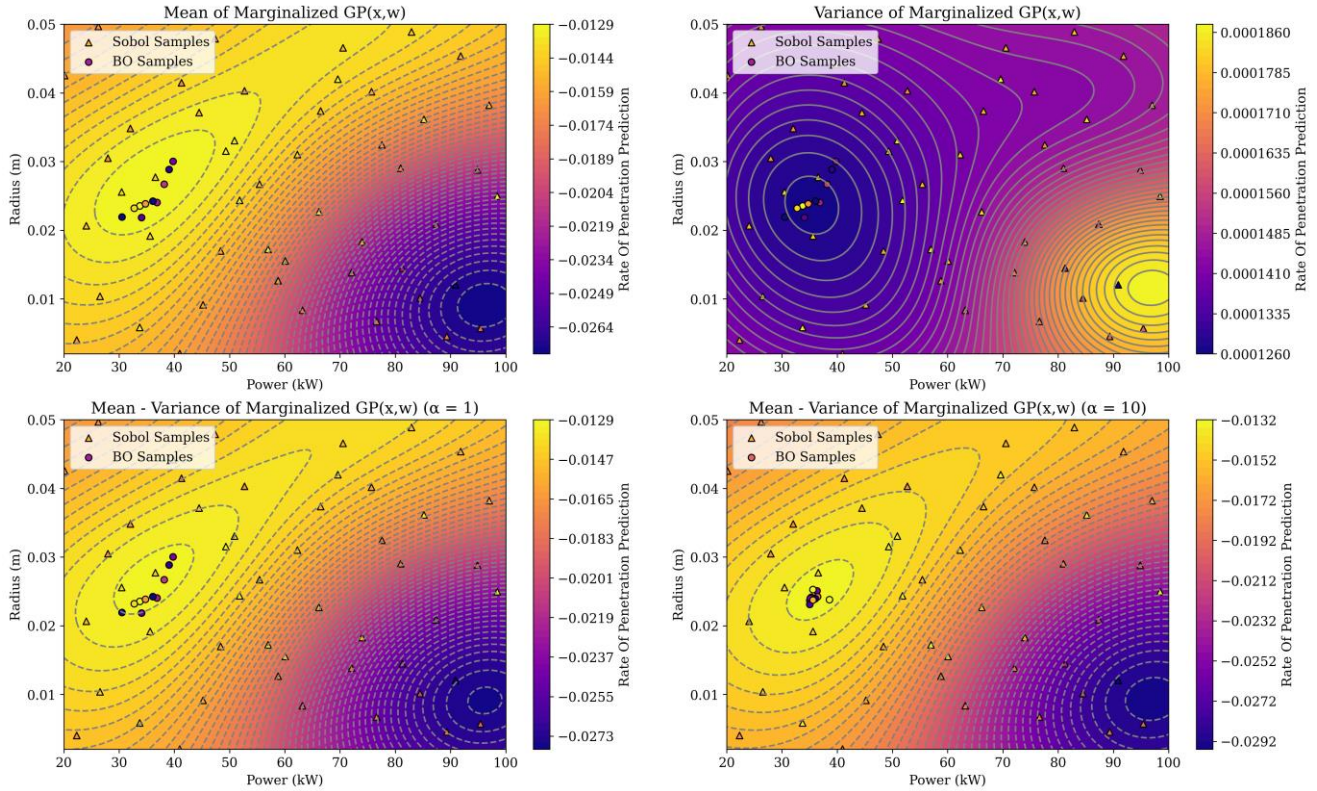


Figure 8: Results of the Mean-Variance Bayesian Optimization run with the Mean-Variance Marginalized Expected Improvement acquisition function. The results of the optimization loop are used to train a GP which makes the predictions for power against beam waist as displayed by the contour lines. The mean and variance predicted by the GP for the $\alpha = 1$ case are shown in the top left and top right figure, while the Mean-Variance results are shown in the bottom left and bottom right for the $\alpha = 1$ and $\alpha = 10$ case respectively.

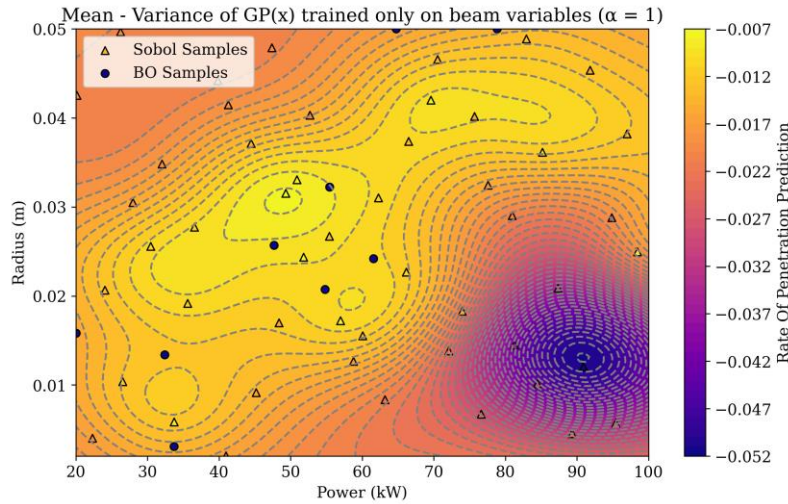


Figure 9: Results of the Mean-Variance Bayesian Optimization run without marginalizing the environmental variables and using the Expected Improvement acquisition function.

The results of the Bayesian optimization algorithm under uncertainty show good convergence for both the $\alpha = 1$ and $\alpha = 10$ case, with the algorithm converging towards the maximum of the Mean-Variance function as predicted by the GP. It can be observed that the location with the highest mean is near the location of the lowest variance, but the two loci do not overlap. As the risk tolerance coefficient increases, we would expect to see the Mean-Variance optima to migrate towards the minima of the variance; the optimal Mean-Variance solution for the $\alpha = 1$ case yielded a beam power of 36.16 kW and radius of 2.62 cm, whereas for the $\alpha = 10$ case, the optima was reached for a beam power of 37.9 kW and radius of 2.53 cm. The benefits of marginalizing the environmental variables are evident by comparing the

results to the case where the GP is trained only on the beam variables and the Expected Improvement acquisition function as shown in Figure 9. The GP struggles to reduce the apparent number of peaks and valleys significantly, inhibiting the convergence of the optimizer.

6. CONCLUSION

In this work, we have described a framework to optimize beam variables for MMW geothermal energy drilling under material uncertainty. We do this by coupling a computational model of the vaporization process with a Bayesian Optimization methodology to sample high performance solutions under a mean-variance objective. The computational model evolves an energy-conservation equation with the MMW beam included as source term, and captures material uncertainty by randomly initializing material grains according to a Voronoi tessellation algorithm and by drawing material properties from a distribution, aiming to reflect variations in granite samples. The Bayesian Optimization algorithm uses a Gaussian Process trained on the beam and environmental variables and the simulation outputs. The environmental variables are then marginalized to obtain mean and variance estimates of the rate of penetration. The acquisition function then computes a marginalized Expected Improvement function and balances it against the output variance to find the next sample location. The algorithm is first applied to a simple test case, with power as tunable variable and latent heat of evaporation as environmental variable, and then against a full high-dimensional test case, with 2 optimizable parameters and several environmental variables. Our results show good performance of the marginalization algorithm and convergence to the optimal beam settings under uncertainty. We have confirmed that the depth reached is sensitive to material properties through our first test, and through our subsequent tests we make the important observation that the optimal beam settings vary depending on the properties of the drilled material and their associated level of uncertainty and gain insights into material uncertainty's relationship to optimal settings.

Future work on the optimization of the MMW ablation technology should aim at considering more complex cases by either adding physics to the model or changing the optimization problem. Adding additional physics into the model can change the optimal beam parameters as the material reacts differently to the beam parameters. The additional physics can include capturing the melt flow of the molten material, the advection of the vaporized material, and the propagation of the electromagnetic wave through air before irradiating the rock. Different optimization problems can vary from considering additional optimization criteria, such as thermal specific efficiency and wellbore cylindricity through a multi-objective criteria, to adding additional optimizable parameters to the study, including the distance of the waveguide from the surface over time or the shape of the incident beam.

REFERENCES

- Alexiades, A. and Autrique, D.: Enthalpy model for heating, melting, and vaporization in laser ablation, *Electronic Journal of Differential Equations Conf*, 19, (2010)
- Ament, S., Daulton, S., Eriksson, D., Balandat, M., & Bakshy, E.: Unexpected improvements to expected improvement for bayesian optimization. *Advances in Neural Information Processing Systems*, 36, (2023), 20577-20612.
- Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., & Bakshy, E.: BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. *Advances in neural information processing systems*, 33, (2020), 21524-21538.
- Bonin, B., Janoušek, V., & Moyen, J. F.: Chemical variation, modal composition and classification of granitoids. *Geological Society, London, Special Publications*, (2020), 491(1), 9-51.
- Bourne, D. P., Pearce, M., & Roper, S. M. : Geometric modelling of polycrystalline materials: Laguerre tessellations and periodic semi-discrete optimal transport. *Mechanics Research Communications*, 127, (2023), 104023.
- Brittenham, T. L., Neudecker, J. W., Rowley, J. C., & Williams, R. E.: Directional drilling equipment and techniques for deep, hot granite wells. *Journal of Petroleum Technology*, 34(07), (1982), 1421-1430.
- Cakmak, S., Astudillo Marban, R., Frazier, P., & Zhou, E.: Bayesian optimization of risk measures. *Advances in Neural Information Processing Systems*, 33, (2020), 20130-20141.
- Cladouhos, T. T., & Callahan, O. A.: Heat extraction from superhot rock: a survey of methods, challenges, and pathways forward. *Transactions—Geothermal Resources Council*, (2023), 2804-51.
- El Amri, R., Le Riche, R., Helbert, C., Blanchet-Scalliet, C., & da Veiga, S.: A sampling criterion for constrained Bayesian optimization with uncertainties. *SMAI Journal of Computational Mathematics*, 9, (2023), 285-309.
- Garnett, R.: *Bayesian optimization*. Cambridge University Press, (2023).
- Girard, A., Rasmussen, C., Candela, J. Q., & Murray-Smith, R.: Gaussian process priors with uncertain inputs application to multiple-step ahead time series forecasting. *Advances in neural information processing systems*, 15, (2020).
- Hamm, S., Hass E., Winick J., Tasca C., Albayrak F., Augustine C., et al. *Geovision: Harnessing the heat beneath our feet*. US Department of Energy, (2019).
- Hong, L. J., Hu, Z., & Liu, G.: Monte Carlo methods for value-at-risk and conditional value-at-risk: a review. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 24(4), (2014), 1-37.
- Houde, M., Araque, C., Oglesby, K., & Woskov, P.: Rewriting the Limits for Deep Geothermal Drilling: Direct Energy Drilling Using Millimeter Wave Technology. In *Proceedings World Geothermal Congress 2020+ 1 Reykjavik*, conference session. (2021).

- Iwazaki, S., Inatsu, Y., & Takeuchi, I.: Mean-variance analysis in Bayesian optimization under uncertainty. In *International Conference on Artificial Intelligence and Statistics* (pp. 973-981), (2021).
- Kuhn, J., Schneider, M., Sonnweber-Ribic, P., & Böhlke, T.: Fast methods for computing centroidal Laguerre tessellations for prescribed volume fractions with applications to microstructure generation of polycrystalline materials. *Computer Methods in Applied Mechanics and Engineering*, 369, (2020), 113175.
- Li, Y. L., Rudner, T. G., & Wilson, A. G. A Study of Bayesian Neural Network Surrogates for Bayesian Optimization. *The Twelfth International Conference on Learning Representations*, (2024).
- Lund, J. W., & Toth, A. N.: Direct utilization of geothermal energy 2020 worldwide review. *Geothermics*, 90, (2021) 101915.
- Makarova, A., Usmanova, I., Bogunovic, I., & Krause, A.: Risk-averse heteroscedastic bayesian optimization. *Advances in Neural Information Processing Systems*, 34, (2021), 17235-17245.
- Mockus, J.: On Bayesian methods for seeking the extremum. In *Proceedings of the IFIP Technical Conference* (pp. 400-404), (1974).
- Oglesby, K., Woskov, P., Einstein, H., & Livesay, B.: Deep geothermal drilling using millimeter wave technology (Final Technical Research Report) (No. DE-EE0005504Final). Impact Technologies LLC, Tulsa, OK (United States) (2014).
- Porlles, J. W., Houde, M. D., Yao, J., Batir, J., Soroush, H., & Madayarov, A.: Borehole Stability Analysis for an Injection Well in Superhot Rock Drilled with Millimeter Wave Technology. In *ARMA US Rock Mechanics/Geomechanics Symposium* (p. D042S059R010), (2024).
- Sani, A., Lazaric, A., & Munos, R.: Risk-aversion in multi-armed bandits. *Advances in neural information processing systems*, 25, (2012).
- Scott, S., Yapparova, A., Weis, P., & Houde, M.: Hydrological constraints on the potential of enhanced geothermal systems in the ductile crust. *Geothermal Energy*, 12(1), (2024), 10.
- Shteinman, A., Anker, Y., & Einat, M. Hard Rock Absorption Measurements in the W-Band. *Journal of Infrared, Millimeter, and Terahertz Waves*, 45(9), (2024), 808-830.
- Wang, L., Niu, X., Chen, S., Liu, J., Liu, Y., Guo, G., & Wang, H.: Design and high-power test of the transmission line for millimeter wave deep drilling. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 33(3), (2020), e2715.
- Williams, C. K., & Rasmussen, C. E. *Gaussian processes for machine learning* (Vol. 2, No. 3, p. 4). Cambridge, MA: MIT press (2006).
- Zhang, A. Z., Millmore, S. T., & Nikiforakis, N.: Thermal simulation of millimetre wave ablation of geological materials. *Computers and Geotechnics*, 161, (2023), 105571.
- Zhang, W., Almgren, A., Beckner, V., Bell, J., Blaschke, J., Chan, C., et al.: AMReX: a framework for block-structured adaptive mesh refinement. *The Journal of Open Source Software*, 4(37), (2019), 1370.