

OPTaaS

OPTaaS is a general purpose global optimizer which is especially designed for:

- Tuning the hyper-parameters of Machine Learning models and Data Science pipelines
- Finding global optima of expensive non-convex objective functions

OPTaaS will identify the optimal parameter configurations more efficiently than alternative approaches because it uses Bayesian Optimization to iteratively learn about the parameter search space which maximizes the efficiency of the process. This is achieved through:

- A *surrogate model* of the underlying optimization problem
- An *acquisition function* which guides the next evaluation points

OPTaaS will recommend configurations to be evaluated on your objective function and will update its surrogate model with each new associated score. This process is sequential and illustrated in Figure 1. A more detailed explanation is provided in our 2D [animation](#).

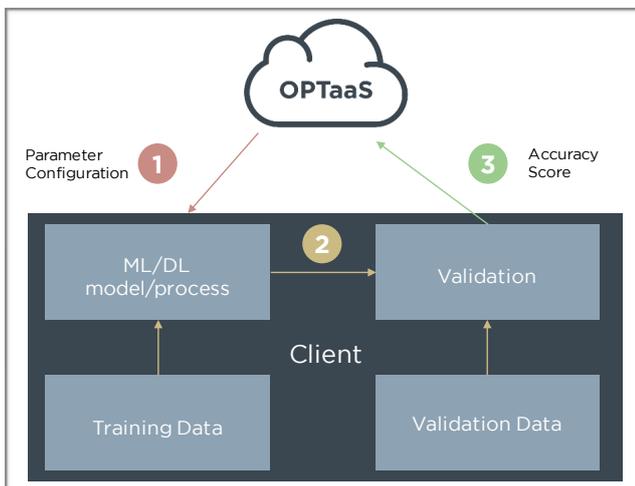


Figure 1: The optimization process with OPTaaS

OPTaaS is deployed via a simple [API](#) which doesn't access your data or models. The optimization process is the following:

1. OPTaaS recommends a configuration to the customer

2. The customer evaluates the configuration on their machines
3. The customer sends back a score (accuracy, Sharpe ratio, Return on investment,...)
4. OPTaaS uses the score to update its surrogate model and the cycle repeats until the optimal configuration has been reached.

Advantages:

- Doesn't access your data or models
- Faster and better results
- Cheaper development
- Increased productivity and creativity
- Interfaces into standard work environments

Technical Features

OPTaaS supports:

- Any type of parameters (continuous, discrete, categorical, conditional, ...)
- Flexible parameter constraints
- Parallelisation of the process through batching
- Simple API client for [Python](#), [Scikit-Learn](#) and [R](#)

Performance

OPTaaS is the best Bayesian Optimizer in the market and has been rigorously benchmarked within a scientific environment against competitive approaches. This environment includes [195](#) well know mathematical functions. The statistical wins (Rate of convergence and accuracy) are summarised in the following tables.

OPTaaS	167
Random	18
Draw	4

OPTaaS	162
HyperOPT	28
Draw	5

OPTaaS	141
SkOPT	43
Draw	11

Additional Features

Batched Bayesian Optimization

OPTaaS can calculate batches of configurations to be evaluated on your cluster. This is particularly interesting when:

- Your objective function is expensive and long to compute
- You have multiple workers available to evaluate configurations

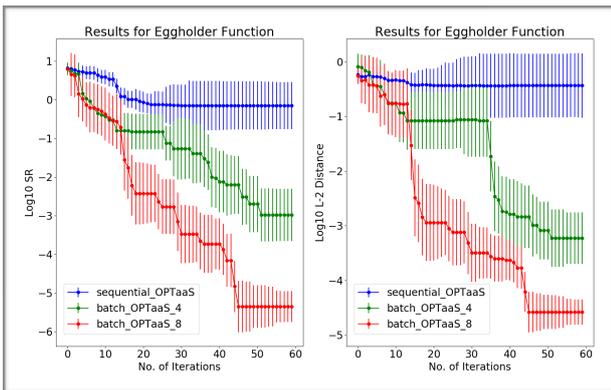


Figure 2: Batched Bayesian Optimization with OPTaaS

OPTaaS models the interactions between the configurations of the batches to optimally sample the search space as well as improve the quality of its recommendations. This significantly reduces clock time.

Multi-Objective Bayesian Optimization

OPTaaS can simultaneously optimize multiple objective functions (e.g. maximize Sharpe, minimize drawdown). In this mode, OPTaaS will attempt to approximate the Pareto frontier.

Meta-Learning and Penalization

OPTaaS can be taught to avoid undesirable configurations by simply penalizing the objective function. For example, in Figure 3, OPTaaS learnt to avoid computationally expensive configurations.

Roadmap

OPTaaS' roadmap has been designed to improve performance and add valuable features:

- **Noisy observations and gradient information:** BO can be improved if gradient information is available but it can also make use of partial and uncertain gradient information.

- **Stopping criteria:** currently BO will stop when it has reached the desired score or iteration budget, however, OPTaaS will be able to intelligently alternate between local and global optimization in order to converge towards the global optimum much faster.
- **BO for dynamic optimization problems:** often the objective function slowly changes over time which requires incorporating these drifts into the target function.

Team and Resources

Mind Foundry is an Oxford University spin-out founded by Professors Stephen Roberts and Michael Osborne who have 35 person years in data analytics. The Mind Foundry team is composed of over 30 world class Machine Learning researchers and elite software engineers, many former post-docs from the University of Oxford. Moreover, Mind Foundry has a privileged access to over 30 Oxford University Machine Learning PhDs through its spin-out status. Mind Foundry is a portfolio company of the University of Oxford and its investors include Oxford Sciences Innovation, the Oxford Technology and Innovations Fund, the University of Oxford Innovation Fund and Parkwalk Advisors.

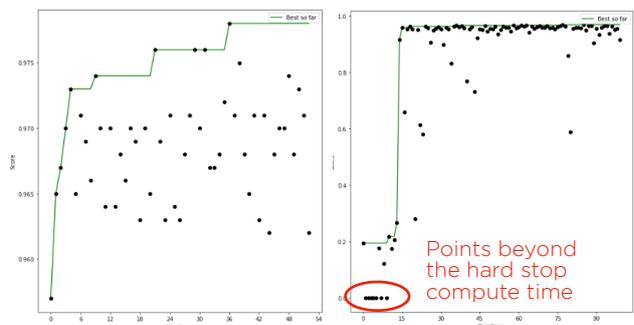


Figure 3: Before and after compute time penalization

Documentation

Tutorials: <https://tutorial.optaas.mindfoundry.ai>

API documentation: <https://optaas.mindfoundry.ai>

Research

<http://www.robots.ox.ac.uk/~mosb/projects/project/2009/01/01/bayesopt/>

Demo: charles.brecque@mindfoundry.ai