

Atlas: A Brain for Self-driving Laboratories

Riley J. Hickman,^{1,2,3,*} Malcolm Sim,^{1,2} Sergio Pablo-García,^{1,2} Ivan Woolhouse,^{1,2} Han Hao,^{1,2} Zeqing Bao,⁴ Pauric Bannigan,⁴ Christine Allen,^{4,5,6} Matteo Aldeghi,^{1,2,3,†} and Alán Aspuru-Guzik^{1,2,3,5,6,7,8,‡}

¹*Chemical Physics Theory Group, Department of Chemistry,
University of Toronto, Toronto, ON M5S 3H6, Canada*

²*Department of Computer Science, University of Toronto, Toronto, ON M5S 3H6, Canada*

³*Vector Institute for Artificial Intelligence, Toronto, ON M5S 1M1, Canada*

⁴*Leslie Dan Faculty of Pharmacy, University of Toronto, Toronto, ON M5S 3M2, Canada*

⁵*Acceleration Consortium, University of Toronto, Toronto, ON M5S 3E5, Canada*

⁶*Department of Chemical Engineering & Applied Chemistry,
University of Toronto, Toronto, ON M5S 3E5, Canada*

⁷*Department of Materials Science & Engineering,
University of Toronto, Toronto, ON M5S 3E4, Canada*

⁸*Lebovic Fellow, Canadian Institute for Advanced Research, Toronto, ON M5G 1Z8, Canada*

(Dated: September 4, 2023)

Self-driving laboratories (SDLs) are next-generation research and development platforms for closed-loop, autonomous experimentation that combine ideas from artificial intelligence, robotics, and high-performance computing. A critical component of SDLs is the decision-making algorithm used to prioritize experiments to be performed. This SDL “brain” often relies on optimization strategies that are guided by machine learning models, such as Bayesian optimization. However, the diversity of hardware constraints and scientific questions being tackled by SDLs require the availability of a set of flexible algorithms that have yet to be implemented in a single software tool. Here, we report ATLAS, an application-agnostic Python library for Bayesian optimization that is specifically tailored to the needs of SDLs. ATLAS provides facile access to state-of-the-art, model-based optimization algorithms—including mixed-parameter, multi-objective, constrained, robust, multi-fidelity, meta-learning, asynchronous, and molecular optimization—as an all-in-one tool that is expected to suit the majority of specialized SDL needs. After a brief description of its core capabilities, we demonstrate ATLAS’ utility by optimizing the oxidation potential of metal complexes with an autonomous electrochemical experimentation platform. We expect ATLAS to expand the breadth of design and discovery problems in the natural sciences that are immediately addressable with SDLs.

Keywords: Self-driving Laboratories, Autonomous Experimentation, Experiment Planning, Bayesian Optimization, Electrochemistry, Python Software, Machine Learning

* riley.hickman@mail.utoronto.ca

† Current address: Bayer Research and Innovation Center, 238 Main St, Cambridge, MA 02142, USA

‡ alan@aspuru.com

I. INTRODUCTION

Self-driving laboratories (SDLs) are advanced technological platforms that use artificial intelligence, robotics, and high-performance computing to perform complex research tasks autonomously, that is, without human intervention. Such platforms aim to streamline and enhance the efficiency of scientific experimentation, research, and analytical processes.^{1–7} SDLs can accelerate the rate at which advanced materials, functional molecules, and industrial processes are designed by enhancing productivity, throughput, accuracy, and reproducibility. Early-stage SDLs have targeted diverse research and development goals, including chemical reaction and process optimization,^{8–18} the design of nanomaterials,^{19–22} and light-harvesting materials,^{23–26} to name a few.^{27–31}

The cornerstone of an SDL is its decision-making algorithm (here informally referred to as its “brain”), which is typically implemented as a data-driven experiment planning strategy. Compared to less dynamic strategies such as Design of Experiment,^{32–34} data-driven approaches leverage feedback from previously completed experiments to inform subsequent recommendations of experimental parameters, resulting in superior sample efficiency. Although many such strategies have been proposed, including gradient-based optimizers,³⁵ evolutionary strategies,^{36–41} and reinforcement learning,^{42,43} Bayesian optimization (BO)^{44–46} has recently emerged as the most popular choice. BO is a sequential optimization strategy for expensive-to-evaluate black-box functions based on machine-learned approximations of the target objective being optimized.

Python libraries for general-purpose BO are plentiful. Popular examples include SCIKIT-LEARN,^{47,48} GPYOPT,⁴⁹ HYPEROPT,^{50–53} SMAC3,^{54,55} DRAGONFLY,⁵⁶ HEBO,⁵⁷ BOTORCH,⁵⁸ AX,⁵⁹ and VIZIER,^{60,61} amongst others. Most of the aforementioned libraries are primarily scoped toward optimization of machine learning (ML) model hyperparameters, and often lack specific functionality requisite for the experimental sciences. For example, the proposed parameters for hyperparameter tuning are generally expected to be executed exactly, while in a SDL it might be impossible to control experimental conditions to high levels of precision.^{62,63} Other requirements for broad applicability of BO in experimental sciences include constrained optimization, with *a priori* known (physical hardware restrictions, safety concerns)^{64,65} and unknown (failed/abandoned synthesis, inadequate conditions for property measurement)^{66–71} constraint functions, as well as the ability for asynchronous experimental execution (*i.e.* recommendation of new parameters before a complete batch of corresponding measurements are available).^{72,73} Although several such libraries do contain the low-level infrastructure necessary to implement more advanced optimization techniques, it remains an expert-level task to correctly organize the required building blocks to produce a working prototype. Software libraries for data-driven decision-making in SDLs and experimental sciences have also been reported.^{19,74–81} While these studies constitute important landmarks in the burgeoning field of SDLs, most target specific experimental frameworks and/or narrow problem types and do not cover the full extent of requirements needed for a truly general-purpose tool.

In this work, we introduce ATLAS, an Object-Oriented Python library for BO that was designed with the broadest applicability to SDLs and experimental science in mind. Figure 1 shows a summary of the main experiment-planning capabilities of ATLAS along with its place within the closed-loop experimentation paradigm. ATLAS intends to provide practitioners of autonomous science with state-of-the-art BO algorithms while abstracting away all complex implementation details. We strived to provide researchers with the numerous, application-agnostic features often required for the successful deployment of BO in practice. This flexibility is expected to allow researchers to focus on customizing their experimental or computational protocols and expand the set of design and discovery problems that

BO and SDLs can tackle.

This paper is organized as follows: Section II provides a brief review of BO and its components. Section III lists the notable features of the ATLAS library and gives code snippet examples for each feature. Finally, Section IV describes a real-world demonstration of ATLAS used in conjunction with CHEMOS 2.0,⁸² an SDL orchestration software, to optimize the oxidation potential of metal complexes in a cyclic voltammetry experiment.

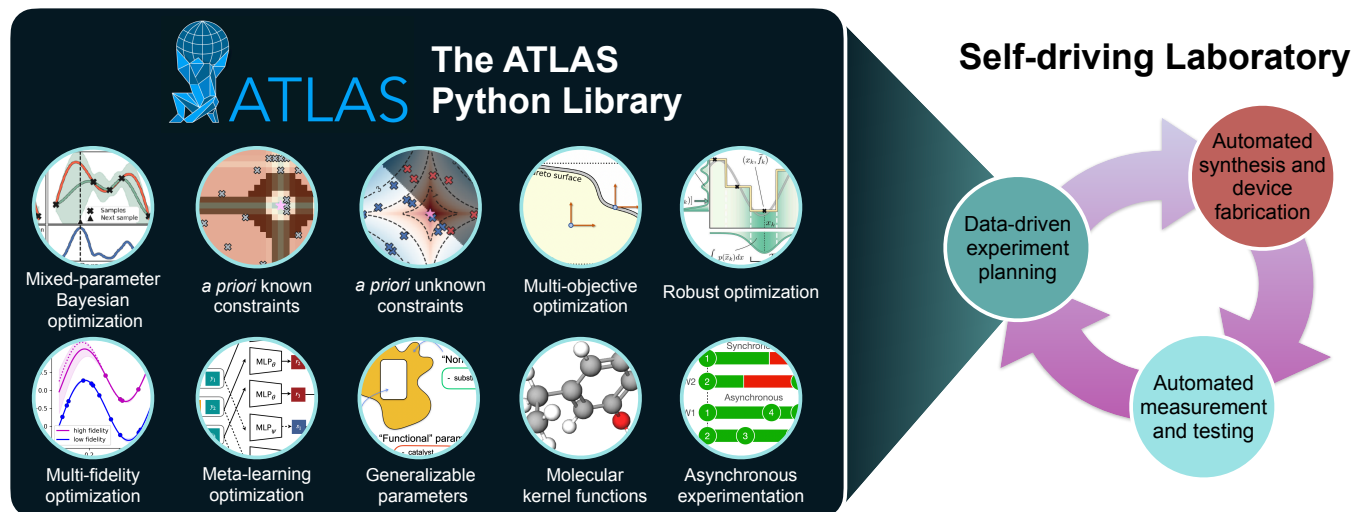


FIG. 1. Conceptual figure showing the important capabilities of ATLAS, as well as its place in a closed-loop experimentation cycle utilized by an SDL.

II. OVERVIEW OF BAYESIAN OPTIMIZATION FOR EXPERIMENT PLANNING

In this section, we briefly review the basic principles of Bayesian optimization as a primer for discussion of the capabilities of ATLAS in Section III.

A. Bayesian optimization

Optimization problems involve identifying parameters, \mathbf{x} , that produce the most desirable outcome for an objective function, $f(\mathbf{x})$. For a minimization problem, the solution is the set of parameters that minimizes $f(\mathbf{x})$,

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (1)$$

where \mathcal{X} is the parameter space, *i.e.* a structured input domain that can be explored during optimization. In a BO setting, the objective function is considered to be a *black-box* function, meaning its structure is *a priori* unknown, and can only be sequentially resolved by empirical measurement. Black-box functions also do not provide access to gradient information, and measurements are typically expected to be corrupted by noise.

Measurements are collected sequentially, either in batches or one-by-one. After collection of a measurement y corresponding to the parameter \mathbf{x} , the surrogate model is trained on the dataset of all available observations, $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$. An acquisition function is then computed based on the surrogate model. The maximum of this function

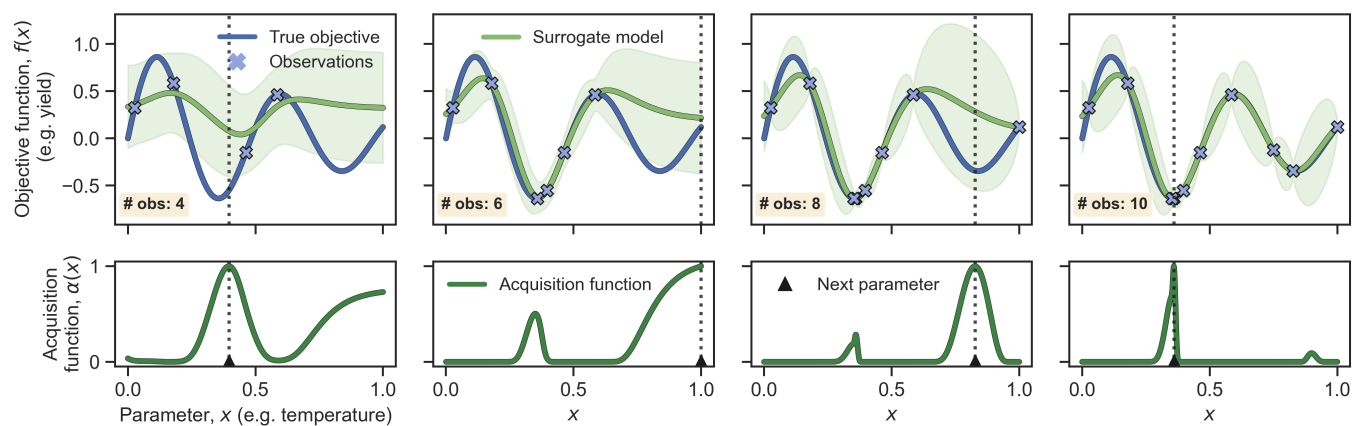


FIG. 2. Conceptual figure showing Bayesian optimization of a 1d function. In this example, the surrogate model is a GP with a Matérn 5/2 kernel, and the acquisition function is the expected improvement criterion.

defines the set of parameters expected to provide maximal utility, and corresponds to the parameter recommended for subsequent measurement. Typically, this iterative procedure is repeated until a pre-defined stopping criterion, such as the exhaustion of an experimental budget, is met. Algorithm 1 shows pseudocode for a BO loop, and Figure 2 visualizes the procedure. The first set of parameters is often not recommended by BO, but are rather produced by random sampling, a low-discrepancy sequence, or an experimental design strategy. This is known as the *initial design phase*.

Algorithm 1: Pseudocode for Bayesian optimization

Data: Parameter space $\mathcal{X} \in \mathbb{R}^d$, objective function $f : \mathcal{X} \mapsto \mathbb{R}$, surrogate model \mathcal{M} , acquisition function $\alpha : \mathcal{X} \mapsto \mathbb{R}$, optimization budget b

Result: Dataset of objective measurements $\mathcal{D}^b = \{(\mathbf{x}_i, y_i)\}_{i=1}^b$

$\mathcal{D}^0 \leftarrow \emptyset$;

$n_{eval} \leftarrow 0$;

while $n_{eval} < b$ **do**

$\mathcal{M} \leftarrow \text{fit } \mathcal{M} \text{ to } \mathcal{D}^{n_{eval}}$;

 /* train surrogate model on current observations */

$\mathbf{x}_{next} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x})$;

 /* optimize acquisition function for next parameter sample */

$y_{next} = f(\mathbf{x}_{next})$;

 /* measure objective function at next parameter sample location */

$\mathcal{D}^{n_{eval}+1} \leftarrow (\mathbf{x}_{next}, y_{next}) \cup \mathcal{D}^{n_{eval}}$;

 /* append newest parameter-measurement pair to dataset */

$n_{eval} \leftarrow n_{eval} + 1$;

1. Acquisition functions

The surrogate model is constructed to approximate the objective function $f(\mathbf{x})$ and can be queried for mean and uncertainty estimates of the objective across the parameter space. Acquisition functions, $\alpha(\mathbf{x})$, are used to guide the selection of parameter recommendations using the surrogate model's prediction. Specifically, maximization of $\alpha(\mathbf{x})$

reveals the parameters \mathbf{x}_{next} for subsequent measurement.

$$\mathbf{x}_{\text{next}} = \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x}). \quad (2)$$

Several acquisition functions are available for use in ATLAS, including the upper and lower confidence bound, expected improvement, probability of improvement, variance-based sampling, and greedy sampling. Our library also features several specialty acquisition functions for more advanced BO concepts, including general parameter optimization (Section III E), multi-fidelity optimization (Section III F), and meta-learning enhanced optimization (Section III G).

2. Acquisition function optimization

A crucial subroutine in BO is the optimization of the acquisition function. Several factors influence the aptitude of optimization strategies for this task, including the types of parameters making up the parameter space and its volume. ATLAS provides 3 distinct acquisition function optimization strategies that are each suited for specific problem types: (i) a constrained gradient optimizer based on SLSQP^{58,83} or ADAM,⁸⁴ (ii) a constrained genetic algorithm (GA) optimizer based on the PYMOO library,⁸⁵ and (iii) a constrained GA based on the DEAP library.^{86,87} Gradient strategies are well-suited for problems with fully-continuous parameter spaces, while GA strategies are well-suited for mixed-parameter problems and fully discrete/categorical problems with large Cartesian product spaces.

B. Gaussian processes

A crucial component of the BO framework is the *surrogate model*, an ML model which produces an estimate of the true objective function given a dataset of observations. Many ML models have been used as BO surrogates, including Bayesian neural networks, tree-based models, and kernel smoothing, but the most prevalent and well-studied choice is the Gaussian process (GP).⁸⁸ GPs are non-parametric probabilistic ML models. They are a collection of random variables such that the joint distribution of each finite set of variables is a multivariate normal. GPs are characterized by a mean, $m(\mathbf{x})$, and covariance function, $k(\mathbf{x}, \mathbf{x}')$, and are written as

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (3)$$

For experiment planning applications, inputs \mathbf{x} are vectors of *parameters* of the experiment. Conveniently, the mean and variance of a GP can be written in closed form. For query parameters, $\hat{\mathbf{x}}$ (*i.e.* those which do not yet have an associated measurement), the GP returns a predictive mean $\hat{\mu}$ and variance $\hat{\sigma}^2$.

1. Kernel functions

The choice of covariance or kernel function $k(\mathbf{x}, \mathbf{x}')$ is an crucial inductive bias for a GP model, and should be selected according to the characteristics of the objective function being modelled. A popular choice for continuous input domains is the Matérn kernel,

$$k(\mathbf{x}, \mathbf{x}') = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{\ell} d(\mathbf{x}, \mathbf{x}') \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\ell} d(\mathbf{x}, \mathbf{x}') \right), \quad (4)$$

where $d(\cdot, \cdot)$ is the Euclidian distance, $K_\nu(\cdot)$ is a modified Bessel function, and $\Gamma(\cdot)$ is the Gamma function. The Matérn kernel also has an additional parameter, ν , which controls the smoothness of the resulting function (fixed at 5/2 in ATLAS), and ℓ is a lengthscale hyperparameter. ATLAS uses a Matérn 5/2 kernel for continuous and numerical discrete input domains, but also supports categorical, mixed continuous/discrete-categorical, and molecular input domains. For categorical spaces in which inputs are one-hot-encoded, we use a kernel function based on Hamming distances,

$$k(\mathbf{x}, \mathbf{x}') = \exp[-d_{\text{Hamming}}(\mathbf{x}, \mathbf{x}')/\ell], \quad (5)$$

where $d_{\text{Hamming}}(\mathbf{x}, \mathbf{x}') = 0$ if $\mathbf{x} = \mathbf{x}'$ and $= 1$ if $\mathbf{x} \neq \mathbf{x}'$. Automatic relevance determination (ARD)⁸⁹ is used for this kernel for all input dimensions. For mixed continuous/discrete-categorical spaces, we use a mixed kernel comprising Matérn 5/2 and Hamming parts,

$$k(\mathbf{x}, \mathbf{x}') = k_{\text{cat}}(\mathbf{x}_{\text{cat}}, \mathbf{x}'_{\text{cat}}) \times k_{\text{cont}}(\mathbf{x}_{\text{cont}}, \mathbf{x}'_{\text{cont}}) + k_{\text{cat}}(\mathbf{x}_{\text{cat}}, \mathbf{x}'_{\text{cat}}) + k_{\text{cont}}(\mathbf{x}_{\text{cont}}, \mathbf{x}'_{\text{cont}}), \quad (6)$$

where \mathbf{x}_{cont} and \mathbf{x}_{cat} are the continuous/discrete and categorical portions of the input vector, respectively. The reader is referred to Section III H for discussion of the kernel used for molecular input domains.

2. Gaussian process training

Fitting a GP to a dataset of observations involves choosing hyperparameters of the kernel function (collectively referred to as $\boldsymbol{\theta}$) and the likelihood noise σ_y^2 . Hyperparameters are chosen by minimizing the negative log marginal likelihood,

$$\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}) = \underbrace{-\frac{1}{2}\mathbf{y}^\top [K_\boldsymbol{\theta}(\mathbf{X}, \mathbf{X}) + \sigma_y^2\mathbf{I}]^{-1}\mathbf{y}}_{\text{promotes data fit}} - \underbrace{\frac{1}{2}\log |K_\boldsymbol{\theta}(\mathbf{X}, \mathbf{X}) + \sigma_y^2\mathbf{I}|}_{\text{penalizes model complexity}} - \frac{N}{2}\log(2\pi). \quad (7)$$

$\mathbf{y} \in \mathbb{R}^n$ is a vector of n objective measurements, $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the *design matrix* of n d -dimensional input vectors. $K_\boldsymbol{\theta}(\mathbf{X}, \mathbf{X})$ is a kernel matrix such that each entry $[K_\boldsymbol{\theta}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. $\sigma_y^2\mathbf{I}$ is the variance of Gaussian noise on the measurements \mathbf{y} . Note that the first term encourages the model's fit to the training data, while the second penalizes overly complex models.⁹⁰ This inherent regularization is attractive for modelling in a low-data setting, such as the initial stages of a sequential model-based optimization campaign.

3. Variational Gaussian process classifier

ATLAS uses a GP-based binary classifier to learn the feasible-infeasible boundaries in parameter space for optimization problems with unknown constraints. This capability is explained in detail in Section III B. Here, we provide details of the GP classifier itself.

For feasibility classification using a GP, exact inference is intractable. Thus, ATLAS approximates the classification posterior using variational inference. Assume a dataset of n binary constraint function $c(\mathbf{x})$ measurements, $\mathcal{D}_c = \{(\mathbf{x}_i, \tilde{y}_i)\}_{i=1}^n$, where $\tilde{y} \in \{0, 1\}$ (0 for feasible measurements and 1 for infeasible measurements). For brevity, we denote the n feasibility measurements as $\tilde{\mathbf{y}} = \{\tilde{y}_i\}_{i=1}^n$ and the design matrix as $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$. GP classification squashes the latent GP output \mathbf{f} through a sigmoidal inverse-link function, $\phi(x) = \int_{-\infty}^x \mathcal{N}(a|0, 1)da$ and a Bernoulli likelihood function conditions the function values. The joint distribution of the feasibility measurements and the latent values is

$$p(\tilde{\mathbf{y}}, \mathbf{f}) = \prod_{i=1}^n \mathcal{B}(\tilde{y}_i | \phi(f_i)) p(\mathbf{f}), \quad (8)$$

where $\mathcal{B}(\tilde{y}_i | \phi(f_i)) = \phi(f_i)^{\tilde{y}_i} (1 - \phi(f_i))^{1 - \tilde{y}_i}$ is the Bernoulli distribution, and $p(\mathbf{f}) = \mathcal{N}(\mathbf{f} | \mathbf{0}, \mathbf{K}_{nn})$ is the usual prior for the values of the GP.

ATLAS' classifier adopts an inducing point approach, in which the latent variables are augmented with additional inducing points. Our strategy follows closely to the one detailed in Hensman *et al.*,⁹¹ where a bound on the marginal likelihood for classification problems is derived. This bound is optimized by adjusting the hyperparameters of the GP kernel, parameters of the multivariate normal variational distribution, and the inducing inputs/points simultaneously using stochastic gradient descent. The classification approach is implemented using the GPyTorch library,⁹² and has the added benefit of scaling more favourably with $|\mathcal{D}_c|$ than does exact GP inference.

III. ATLAS LIBRARY OVERVIEW

ATLAS users interact with the package *via* a high-level “ask-tell” interface, in which a `Planner` instance is iteratively queried for parameters, some physical or computational measurement is completed, and the parameter-measurement pairs are added to a `OLYMPUS Campaign` instance. `OLYMPUS`^{79,80} is a complementary experiment-planning framework developed by our research group that provides an interface to ATLAS. For example, `OLYMPUS` implements an abstraction for a generic experiment planning strategy, from which all ATLAS strategies inherit. Also, the `OLYMPUS Campaign` object is used for storing optimization trajectories and meta-information. `OLYMPUS` also provides definitions of parameter types and spaces, and achievement scalarizing functions for multi-objective optimization, all of which are used by ATLAS.

To demonstrate the usage of our software, we present a minimal code example in which the `GPPlanner` (Gaussian Process Planner) from ATLAS is used to minimize the Branin-Hoo surface,⁹³ $f : \mathcal{X} \in \mathbb{R}^2 \mapsto \mathbb{R}$. “Ask-tell” experimentation proceeds iteratively by generating parameters to be measured using the planner's `recommend()` method, and informing the `OLYMPUS Campaign` instance about the corresponding measurement using its `add_observation()` method. We opt to use a flexible “ask-tell” interface to remain application-agnostic, as well as allow for analysis and customization of the optimization loops. Measurement steps usually involve calls to specialized robotic laboratory equipment or computational simulation packages, which can be fully customized by the user. Figure 3 visualizes the results of this simple example. The `num_init_design` argument to the `GPPlanner` constructor defines the number of parameters to recommend in the initial design phase. This value defaults to 5 and will not be listed as an argument in subsequent examples for brevity.

```

from olympus import Surface, Campaign
from atlas.planners.gp.planner import GPPlanner

surface = Surface(kind='Branin') # instantiate 2d Branin-Hoo objective function

campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(surface.param_space) # save details of the optimization domain into the campaign object

planner = GPPlanner(goal='minimize', num_init_design=5) # instantiate Atlas planner
planner.set_param_space(surface.param_space) # define the optimization domain for the planner

while campaign.num_obs < 30: # run no more than 30 experiments
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        measurement = surface.run(sample) # measure Branin-Hoo function
        campaign.add_observation(sample, measurement) # tell planner about most recent observation

```




FIG. 3. Visualization of the minimal code optimization of the 2d Branin-Hoo surface using ATLAS. The left-most subplot shows a contour plot of the Branin-Hoo surface with its triply-degenerate global minimum highlighted with pink stars. The center subplot shows the location of the parameters recommended by ATLAS as gray crosses. The right-most subplot shows the optimization regret trace.

ATLAS supports parameter spaces consisting of continuous, discrete, and categorical parameters, and arbitrary combinations thereof, in sequential or batched optimization mode. The definition of vector-valued descriptors for categorical parameter options is also supported.⁷⁸ BO strategies in ATLAS primarily use GP surrogate models,⁸⁸ and use the low-level infrastructure of the PyTorch,⁹⁴ GPyTorch⁹² and BoTorch⁵⁸ libraries. In the following section, the main capabilities of ATLAS (Figure 1) are explained. For more information on each concept, please visit the ATLAS GitHub,⁹⁵ documentation, and tutorial notebook. Importantly, ATLAS allows users to combine its key capabilities to suit their specialized needs. Barring a few incompatibilities (described in detail in our documentation), capabilities can

be combined and interchanged freely. For example, the *a priori* known and unknown constraints can be used for any parameter space, with any acquisition function, acquisition optimizer, and any planner. Multi-objective optimization via ASFs, robust optimization, and asynchronous optimization can also be used in this way.

A. *A priori* known constraints

Using a simple interface, users can specify arbitrary known constraint functions on the parameter space, which results in portions of the space being omitted from consideration by planners.⁶⁴ We also supply convenient ways to specify commonly occurring known constraint types, including compositional (simplex),⁶⁵ permutation (ordering),⁶⁵ pending experiments, or process-constrained batches.⁹⁶

The following code snippet shows the instantiation of the `GPPlanner` with a user-defined constraint function for the Dejong surface, $f : \mathcal{X} \in \mathbb{R}^2 \mapsto \mathbb{R}$. Constraint functions are Python callables which return a boolean value, `True` (`False`) for feasible (infeasible) parameters, and are passed to the constructor of `GPPlanner` as a list using the `known_constraints` argument.

```

from olympus import Surface, Campaign
from atlas.planners import GPPlanner

# define the known constraint function
def known_constraint(params):
    # params is a array-like object representing one parameter setting
    y = (params[0]-0.5)**2 + (params[1]-0.5)**2
    if 0.05 < y < 0.15:
        return False
    return True

surface = Surface(kind='Dejong') # instantiate 2d Branin-Hoo objective function
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(surface.param_space)

planner = GPPlanner(goal='minimize', known_constraints=[known_constraint])
planner.set_param_space(surface.param_space)

```

Next, we show an example of a special known constraint case built into ATLAS: compositional or simplex constraints. This constraint type is useful when parameters (or a subset thereof) must lie on a standard n -simplex, *i.e.* $\Delta^n = \{(x_0, \dots, x_n) \in [0, 1]^n \mid \sum_{i=0}^n x_i = 1, x_i \geq 0 \forall i\}$. Such constraints are commonly encountered in SDL applications.⁶⁵

To demonstrate this constraint type, we use the `oer_plate_a` dataset from OLYMPUS, which reports high-throughput screens for oxygen evolution reaction (OER) activity by systematically exploring high-dimensional chemical spaces.^{97,98} The dataset is a discrete library of 2121 catalysts, comprising all unary, binary, ternary and quaternary compositions from unique 6 element sets at 10% intervals. The composition system of the

oer_plate_a dataset is Mn-Fe-Co-Ni-La-Ce. OLYMPUS emulates the discrete dataset with a Bayesian neural network (BNN) to produce noisy virtual measurements.⁹⁹ Fractional compositions must lie on the standard 6-simplex $\Delta^6 = \{(x_0, \dots, x_5) \in [0, 1]^6 \mid \sum_{i=0}^5 x_i = 1, x_i \geq 0 \forall i\}$. The following code snippet shows instantiation of the GPPlanner for this problem. The compositional_params argument takes a list of integers representing the parameter space indices to be treated with the compositional constraint. In this case, all 6 parameters are subject to the constraint.

```

from olympus.emulators import Emulator
from atlas.planners import GPPlanner

# instantiate the BNN emulator for the `oer_plate_a` dataset
emulator = Emulator(dataset='oer_plate_a', model='BayesNeuralNet')

campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(emulator.param_space)

planner = GPPlanner(goal='minimize', compositional_params[0,1,2,3,4,5])
planner.set_param_space(emulator.param_space)

```

Lastly, we show an example using the pending experiment known constraint type. The interpretation of this constraint is simple: parameters that have been assigned to measurement, but for which the experiments have not been completed yet, must be avoided by the planner to avoid duplicate recommendation. ATLAS provides a simple method for all planners to set pending experiments. The planner's set_pending_experiments() method can be called at any time within an optimization campaign to inform the planner about parameter settings to be avoided. This method takes as an argument pending_experiments, which is a list of OLYMPUS ParameterVector objects. Each subsequent call overwrites the pending parameters from the last iteration. To remove all the pending experiments, one can use the planner's remove_pending_experiments() method.

```

from olympus import Surface, Campaign
from atlas.planners import GPPlanner

surface = Surface(kind='CatCamel') # instantiate 2d categorical Camel surface

campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(surface.param_space)

planner = GPPlanner(goal='minimize') # instantiate Atlas planner
planner.set_param_space(surface.param_space)

while campaign.num_obs < 30:
    pending_experiments = get_my_pending_exps() # user defined pending experiments

```

```

planner.set_pending_experiments(pending_experiments) # inform planner about pending experiment parameters
samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
for sample in samples:
    measurement = surface.run(sample) # measure CatCamel function
    campaign.add_observation(sample, measurement) # tell planner about most recent observation

```

B. *A priori* unknown constraints

The inclusion of NaN (not a number) objective values is supported, which could occur due to attempted but failed experimental measurements. Several strategies are provided which learn the unknown constraint function on the fly, using a GP-based binary feasibility classifier (explained in detail in Section II B 3). These predictions are used in conjunction with the typical regression surrogate model to parameterize feasibility-aware acquisition functions, $\alpha_c(\mathbf{x})$. All acquisition function types in ATLAS have a feasibility-aware analogue.

We provide an example optimization of the Branin-Hoo surface with an *a priori* unknown constraint function $c(\mathbf{x})$, visualized by the shaded region in Figure 4. The example uses the feasibility-interpolated acquisition strategy (fia-1) with the UCB acquisition function. Two additional arguments to the GPPlanner constructor are required for optimization with unknown constraints. `feas_strategy` indicates the feasibility-aware acquisition type, and the `feas_param` argument indicates the associated parameter. The fia strategies interpolate between the vanilla acquisition function $\alpha(\mathbf{x})$ and the conditional output of the feasibility classifier, $P(\text{feasible}|\mathbf{x})$ using the following expression:

$$\alpha_c(\mathbf{x}) = (1 - c^t) \times \alpha(\mathbf{x}) + c^t \times P(\text{feasible}|\mathbf{x}). \quad (9)$$

c is the ratio of infeasible measurements to total measurements, and $t \in \mathbb{R}_+$ is a parameter (specified with the `feas_param` argument) which controls risk when it comes to selecting infeasible parameters. Here, smaller values of t de-emphasize the feasibility classifier's contribution (second term in Equation 9) and thus indicate more risk, while larger values do the opposite and represent less risk.

```

from olympus import Campaign
from atlas.planners import GPPlanner

# instantiate 2d constrained Branin-Hoo objective function (available on GitHub repo)
surface = BraninConstr()
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(surface.param_space)

planner = GPPlanner(
    goal='minimize',
    feas_strategy='fia',
    feas_param=1.,
    acquisition_type='ucb',

```

```

) # instantiate Atlas planner
planner.set_param_space(surface.param_space)

while campaign.num_obs < 100:
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        measurement = surface.run(sample) # measure constrained Branin-Hoo function
        campaign.add_observation(sample, measurement) # tell planner about most recent observation

```

Figure 4 shows the results of a larger benchmark of feasibility-aware acquisition strategies in ATLAS on the 2d constrained Branin-Hoo surface. The legend of the center subplot lists the unknown constraint strategies and associated parameters available for use in ATLAS. We omit a full discussion and benchmark of these strategies, as they will be thoroughly detailed in an upcoming manuscript. Briefly, the `naive-0` strategy is a simple baseline approach which does not use the feasibility classifier. Instead, the NaN objective value of infeasible measurements is replaced by the current worst feasible measurement in \mathcal{D} .⁷⁰ Although this strategy is effective for avoiding infeasible measurements (lowest % infeasible measurements in right-most subplot), optimization performance is sacrificed, especially when the optimum of the problem is located close to an infeasible region. `fia-1` is among the top performers for this example.




FIG. 4. Visualization of an *a priori* unknown constraint optimization benchmark on the 2d Branin-Hoo surface using ATLAS. The left-most subplot shows a contour plot of the Branin-Hoo surface with its triply-degenerate global minimum highlighted with pink stars, and the constrained regions shaded. The plot also shows the locations of the (in)feasible parameters recommended by ATLAS as (white)gray crosses (recommended using the `fia-1` strategy). The center subplot shows regret traces for each strategy averaged over 100 independent executions. The right-most subplot shows distributions of percentages of infeasible measurements (*i.e.* NaN objective values) produced by each strategy during a run.

C. Multi-objective optimization

ATLAS supports multi-objective optimization for all planners by using achievement scalarizing functions (ASFs) defined in OLYMPUS. Multi-objective problems feature an objective space $\mathcal{Y} \in \mathbb{R}^n$ corresponding to a set of $n > 1$

objective functions $\mathbf{f} = \{f_i\}_{i=1}^n : \mathcal{X} \mapsto \mathcal{Y}$ to be optimized concurrently. ASF \mathcal{S} transforms a vector of objective measurements to a scalar merit value, $\mathcal{S} : \mathcal{Y} \mapsto [0, 1]$ which is processed by the optimizer. Available ASF types are CHIMERA,¹⁰⁰ HYPERVOLUME,^{101–104} CHEBYSHEV,^{105,106} and WEIGHTED SUM.^{107,108}

As an illustrative example of multi-objective optimization in ATLAS, we use the `dye_lasers` dataset from OLYMPUS, which reports computed photophysical properties for 3458 organic molecules synthesized from three groups of molecular building blocks – A, B, and C (resulting in A-B-C-B-A pentamers shown in Figure 5a).¹⁰⁹ Each molecule was subjected to a computational protocol consisting of cheminformatic, semi-empirical and *ab initio* quantum chemical steps to compute absorption and emission spectra, as well as fluorescence rates. The objectives for this dataset are i) the peak score, a dimensionless quantity given by the fraction of the fluorescence power spectral density that falls within the 400 - 460 nm region (to be maximized), ii) the spectral overlap of the absorption and emission spectra (to be minimized), and iii) the fluorescence rate (to be maximized).

The HYPERVOLUME ASF is used for this example. One must include additional arguments to the `GPPlanner` constructor, namely, a boolean value for `is_moo`, the name of the ASF for `scalarizer_kind`, the objective space \mathcal{Y} as an OLYMPUS `ParameterSpace` object, and a list of `goals` representing the individual optimization goals for each objective (either "max" or "min"). For parameter spaces with categorical parameters, note that we can toggle between using a descriptor representation for the options and one-hot encodings using the `use_descriptors` argument.

```

from olympus import Dataset, Campaign
from atlas.planners.gp.planner import GPPlanner

dataset = Dataset(kind='dye_lasers') # instantiate the `dye_lasers` dataset from Olympus
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(dataset.param_space)

planner = GPPlanner(
    goal='minimize', # overall goal must always be set to `minimize` for moo problems
    use_descriptors=False,
    is_moo=True,
    scalarizer_kind='Hypervolume',
    value_space=dataset.value_space,
    goals=['max', 'min', 'max'], # individual goals for each objective
) # instantiate Atlas planner
planner.set_param_space(dataset.param_space)

while campaign.num_obs < 200:
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        measurement, _, __ = dataset.run(sample) # measure objectives of `dye_laser` dataset
        campaign.add_observation(sample, measurement) # tell planner about most recent observation

```

Figure 5b shows the results of a larger scale benchmark experiment where the performance of the `GPPlanner` is

compared to a random search on the `dye_lasers` dataset. Each subplot shows traces of the objective values associated with the measurement assigned the best hypervolume as a function of the number of evaluations. While the peak score and fluorescence rate traces are comparable between strategies, the GPPlanner is able to identify candidate molecules with lower spectral overlap between absorption and emission spectra (a proxy for reduced losses from self-absorption of emitted light) significantly quicker than random sampling. This indicates better multi-objective optimization performance.

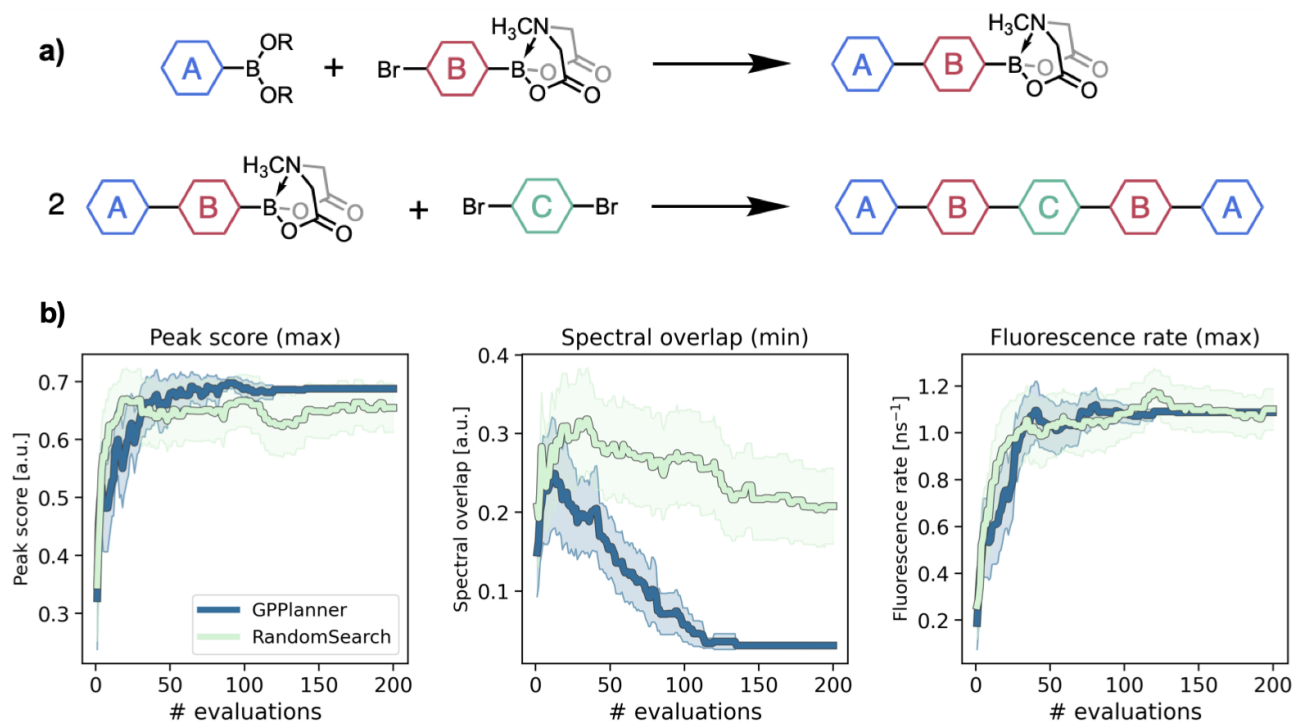


FIG. 5. a) Reaction scheme of iterative Suzuki-Miyaura cross-coupling reaction proposed to synthesize symmetric A-B-C-B-A pentamers in the `dye_lasers` dataset. b) Individual objective traces corresponding to the experiment with the best hypervolume value as a function of the number of transpired experiments. Solid lines represent the mean objective values averaged over 50 independent runs and shaded regions represent the 95% confidence interval.

D. Robust optimization

For all planners and parameter types, we have integrated the GOLEM algorithm,⁶² which allows users to identify optimal solutions that are robust to input parameter uncertainty. This helps ensure reproducible performance of optimized experimental protocols and processes.

In order to demonstrate how GOLEM is integrated into ATLAS, we reproduce the setup of the noisy high-performance liquid chromatography (HPLC) protocol optimization experiment from the original publication.⁶² In this application, an HPLC protocol is calibrated by adjusting 6 process parameters with the goal of maximizing the amount of drawn sample reaching the detector (referred to as the *peak area*). It is assumed that input parameters P1 (`sample_loop`) and P3 (`tubing_volume`) are subject to significant noise, and that the other four parameters are noiseless. Normally

distributed noise truncated at zero is used for both parameters, with standard deviations of 0.008 mL and 0.08 mL for P1 and P3, respectively. The HPLC response is emulated within the OLYMPUS package using BNNs (hplc dataset).

The problem setup for this example is shown in the following code snippet. For all planners, the constructor argument `golem_config` is available, which expects a dictionary whose keys are parameter names for which the user would like to specify an input uncertainty distribution. The corresponding values are themselves dictionaries, for which the user must specify the distribution type and its associated parameters. GOLEM ships with a diverse set of distribution types, which are detailed in its documentation. For all parameters *not* itemized within the `golem_config` argument, ATLAS automatically assigns them a Delta distribution, meaning no uncertainty/noiseless.

```

from olympus import Emulator, Campaign
from atlas.planners.gp.planner import GPPlanner

# instantiate the `hplc` experiment emulator from Olympus
emulator = (dataset='hplc', model='BayesNeuralNet')
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(emulator.param_space)

planner = GPPlanner(
    goal='maximize',
    golem_config={
        'sample_loop': {
            "dist_type": "TruncatedNormal",
            "dist_params": {"std": 0.008, "low_bound": 0.0},
        }, # P1 distribution config
        'tubing_volume': {
            "dist_type": "TruncatedNormal",
            "dist_params": {"std": 0.08, "low_bound": 0.0 },
        } # P3 distribution config
    }
) # instantiate Atlas planner
planner.set_param_space(emulator.param_space)

while campaign.num_obs < 50:
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        measurement, _, __ = emulator.run(sample) # measure `hplc` response
        # tell planner about most recent observation
        campaign.add_observation(sample, measurement)

```

E. Optimization for generalizable parameters

ATLAS includes a strategy based on BO and variance-based active learning¹¹⁰ to identify sets of parameters from \mathcal{X} that result in average-best performance over a set of variables $S = \{s_i\}_{i=1}^N$. Instead of optimizing objective function $f(\mathbf{x})$, the objective $f(\mathbf{x}, s) = \frac{1}{N} \sum_{i=1}^N \tilde{f}(\mathbf{x}, s_i)$ is targeted. Our strategy is inspired by the approach reported by Angello *et al.*,¹¹¹ which was used to design chemical reaction conditions resulting in high yields across a range of substrates. Importantly, this approach alleviates one from having to measure the full objective function for each recommended set of parameters, which can become costly when the number of general parameter options are large.

As an illustrative example, we consider the `suzuki_i-suzuki_iv` datasets from OLYMPUS, which report the yield and catalyst turnover number for flow-based Suzuki-Miyaura coupling reactions with varying substrates shown in Figure 6a and b.¹¹² There are three continuous parameters (temperature, residence time, and catalyst loading) and one categorical parameter (Pd catalyst ligand). The objective is to simultaneously maximize both the yield and catalyst turnover number across all four substrates with the general parameter optimization strategy. Since this is also a multi-objective optimization problem, we use the HYPERVOLUME ASF.

The following code snippet describes the setup for this example. We add an additional categorical parameter named "s", which comprises the general parameter options (in our case these are the different possible substrates for the Suzuki-Miyaura reaction encoded as Roman numerals). For general parameter problems, the `acquisition_type` argument of the planner must be set to "general". Similar to the known constraints problems outlined in Section III A, the constructor takes an argument called `general_parameters`, which must be a list of integers representing the parameter space indices of those parameters to be treated as general parameters. In our case, only the first parameter, "s", is a general parameter.

```

from olympus import Emulator, Campaign
from atlas.planners.gp.planner import GPPlanner

# instantiate the `suzuki_i`-`suzuki_iv` BNN emulators from Olympus
emulators = {
    'i': Emulator(dataset='suzuki_i', model='BayesNeuralNet'),
    'ii': Emulator(dataset='suzuki_ii', model='BayesNeuralNet'),
    'iii': Emulator(dataset='suzuki_iii', model='BayesNeuralNet'),
    'iv': Emulator(dataset='suzuki_iv', model='BayesNeuralNet'),
}

# define measurement function
def measure(func_params, s):
    measurement, _, __ = emulators[s].run(func_params)
    return measurement

# create parameter space with general parameter
param_space = ParameterSpace()
param_space.add(ParameterCategorical(name='s', options=['i', 'ii', 'iii', 'iv'])) # general parameter

```



```

for param in emulator_i.param_space: # functional parameters
    param_space.add(param)

campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(param_space)

planner = GPPlanner(
    goal='minimize', # overall goal must always be set to `minimize` for moo problems
    is_moo=True,
    scalarizer_kind='Hypervolume',
    value_space=emulator_i.value_space,
    goals=['max', 'max'],
    acquisition_type='general', # acquisition function type must be 'general'
    general_parameters=[0], # list of indices of general parameters in parameter space
) # instantiate Atlas planner
planner.set_param_space(param_space)

while campaign.num_obs < 50:
    samples = planner.recommend(campaign.observations)
    for sample in samples:
        measurement, _, __ = measure(sample, sample.s)
        # tell planner about most recent observation
        campaign.add_observation(sample, measurement)

```

Figure 6c shows the results of a larger scale benchmark experiment on this problem. We compare the performance of the general parameter optimization strategy in ATLAS to a strategy in which, for each set of recommended parameters, we measure the *full objective*, i.e. $f(\mathbf{x}, s) = \frac{1}{4} \sum_{i=1}^4 \tilde{f}(\mathbf{x}, s_i)$. Effectively, the latter strategy must make 4 objective measurements for each set of recommended parameters, while the general parameter strategy must only make 1. The box plots in Figure 6c show the hypervolume of solutions identified by each strategy averaged over the 4 reaction types. On average, the general parameter strategy produces larger hypervolumes, indicating superior multi-objective optimization performance.

F. Multi-fidelity optimization

Multi-fidelity BO targets problems where two or more “information sources” are available to the researcher. Typically, the information sources generate measurements of the same property at different levels of fidelity, precision, or accuracy, and are available at varying cost. For instance, a chemical property could be estimated using a crude but inexpensive simulation (low-fidelity) as a proxy for an accurate but expensive experimental determination (high-fidelity). Multi-fidelity strategies are quickly becoming a popular approach for resource-intensive problems in chemistry and materials science.^{113–119} ATLAS provides a `MultiFidelityPlanner` based on the trace-aware knowledge gradient^{120,121} and augmented-EI (aEI) acquisition functions¹²² which allows for the inclusion of an arbitrary number of information




FIG. 6. a) Scheme for the Suzuki-Miyaura cross-coupling of two heterocycles in the presence of 1,8-diazabicyclo[5.4.0]undec-7-ene (DBU) and THF/water. b) Structure of the substrates in each of the four Suzuki-Miyaura reaction cases, corresponding to the `suzuki_{i,ii,iii,iv}` datasets and general parameter. c) Results of the comparative optimization experiment. Boxplots show the normalized average hypervolume across the four reactions for solutions from 50 independent runs (larger hypervolume is better). The experimental budget for these runs was 30 measurements.

sources with discrete fidelity levels.

To illustrate multi-fidelity BO with ATLAS, we use a dataset of simulated band gaps for 192 hybrid organic-inorganic perovskite (HOIP) materials reported by Kim *et al.*¹²³ HOIP candidates are designed from a set of 4 halide anions, 3 group-IV cations and 16 organic anions. Electronic and geometric descriptors of the HOIP components are available through OLYMPUS. Two density functional theory (DFT) information sources are available for all 192 HOIP candidates.

- *low-fidelity*: Band gaps computed using the generalized gradient approximation (GGA), E_g^{GGA} .¹²⁴
- *high-fidelity*: Band gaps computed using the Heyd-Scuseria-Ernzerhof (HSE06) exchange-correlation functional, E_g^{HSE06} .^{125,126}

The GGA level of theory is computationally feasible for these systems but is known to underestimate E_g by 30%.¹²⁴ The HSE06 level of theory is expected to be on par with experimentally determined band gaps but is computationally restrictive. We omit a detailed comparison of the computational methods and estimate the HSE06 level of theory to be an order of magnitude more expensive than the GGA level.

The following code snippet sets up this multi-fidelity optimization problem using ATLAS' `MultiFidelityPlanner`. We've defined helper functions to measure the objective at each fidelity level (`lookup` table provided on GitHub repo), and a function to compute the cumulative experimental cost. The `ParameterSpace` is constructed with an additional fidelity parameter, "s", which is a `ParameterDiscrete` instance with options corresponding to the expense of low-fidelity information sources relative to the high- or target fidelity source organized in increasing order. By convention, ATLAS expects the target fidelity to have a value of 1.0, while the lower fidelity levels have values $0.0 < s < 1.0$. Here, the choice of 0.1 for the E_g^{GGA} determinations reflects our estimate that GGA is an order of magnitude cheaper than

HSE06. The constructor of the `MultiFidelityPlanner` must receive one additional argument, `fidelity_params`, which is the parameter space index of the fidelity parameter "s".

```

from olympus import Dataset, Campaign
from atlas.planners import MultiFidelityPlanner

dataset = Dataset(kind='perovskites') # instantiate `perovskites` dataset from Olympus
# build parameter space with fidelity parameter
param_space = ParameterSpace()
param_space.add(ParameterDiscrete(name='s',options=[0.1, 1.0])) # fidelity parameter
for param in dataset.param_space: # add perovskite component parameters ('organic', 'cation', and 'anion')
    param_space.add(param)

campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(param_space)

# instantiate Atlas planner
planner = MultiFidelityPlanner(goal='minimize', use_descriptors=True, fidelity_params=0)
planner.set_param_space(param_space)

cumul_cost = 0.
while cumul_cost < 50.:
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        # measure the HOIP band gap at one of two fidelity levels (user defined)
        measurement = measure(sample, fidelity=sample.s)
        campaign.add_observation(sample, measurement) # tell planner about most recent observation
        # compute cumulative experiment cost (user defined)
        cumul_cost = compute_cost(campaign.observations.get_params())

```

Note that for this example, we define the BO stopping criterion to be a cumulative experimental cost budget rather than the number of transpired objective evaluations. By default, the `MultiFidelityPlanner` automatically determines which fidelity level to measure the objective at each iteration. However, an SDL researcher may want to further customize their multi-fidelity optimization campaign such that, for example, they can alternate between batches of low- and high-fidelity measurements. ATLAS enables such customized campaigns by allowing the user to specify the fidelity level they wish for the parameters to be measured for the upcoming batch of recommendations. One may set the `MultiFidelityPlanner`'s `current_ask_fidelity` attribute by calling the `set_ask_fidelity` method and specifying the desired level. ATLAS employs constrained acquisition function optimization to deliver the desired parameter recommendations. The following code snippet revisits the HOIP example and assumes the researcher desires to alternate between low- and high-fidelity measurements.

```

cumul_cost = 0.
iter_ = 0
while cumul_cost < 50:
    if iter_ % 2 == 0:
        planner.set_ask_fidelity(1.0) # measure target fidelity on even iterations
    else:
        planner.set_ask_fidelity(0.1) # measure low-fidelity on odd iterations
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:
        measurement = measure(sample) # measure the HOIP bandgap at one of two fidelity levels
        campaign.add_observation(sample, measurement) # tell planner about most recent observation
        cumul_cost = compute_cost(campaign.observations.get_params()) # compute cumulative experiment cost
    iter_ += 1

```

Figure 7 shows the results of a larger-scale benchmark comparison between ATLAS’ `MultiFidelityPlanner` and `GPPlanner` on the perovskites problem. Boxplots show the cumulative simulation cost taken to identify the perovskite with the lowest HSE06 bandgap, E_g^{HSE06} in the dataset over 100 independently seeded runs for each strategy. The `MultiFidelityPlanner` also has access to bandgap measurements at the GGA level (with an E_g^{GGA} to E_g^{HSE06} measurement ratio of 8:1), while the `GPPlanner` has access to E_g^{HSE06} measurements only. The ratio of low- to high-fidelity measurements is set to 8:1 for the `MultiFidelityPlanner`. On average, the `GPPlanner` achieves the lowest high-fidelity bandgap with a cost of 13.9 ± 0.56 a.u., while the `MultiFidelityPlanner` achieves this with a cost of only 9.6 ± 0.35 a.u. These results demonstrate the ability of multi-fidelity BO strategies to leverage inexpensive measurements to augment cost-effective optimization of resource-intensive objectives in SDLs.

G. Meta-/few-shot learning enhanced optimization

With ATLAS, users may easily incorporate data from historical optimization campaigns. These source tasks can be leveraged to accelerate the optimization rate on a novel campaign by using one of two meta-/few-shot learning planners: the Ranking-Weighted Gaussian Process Ensemble planner (`RGPEPlanner`)^{127,128} and the Deep Kernel Transfer planner (`DKTPlanner`).^{129,130} Importantly, these strategies can each transcend the innate design restrictions of typical BO by inferring an inductive bias implicitly from data. In SDLs, this amounts to learning inductive biases that closely resemble particular concepts in chemistry or materials science, and then applying them to related optimization problems. Such approaches have been used to optimize chemical reactions in SDL applications.^{131,132}

We use the `buchwald_{a,b,c,d,e}` datasets from OLYMPUS to showcase the aptitude of meta-/few-shot learning planners to accelerate optimization given historical optimization campaign data. The `buchwald` datasets comprise 5 datasets which each report the yield of Pd-catalyzed Buchwald-Hartwig amination reactions of aryl halides (3 options) with 4-methylaniline in the presence of varying isoxazole additives (22 options), Pd catalyst ligands (4 options), and bases (3 options).¹³³ Each dataset consists of 792 yield measurements. The reaction scheme is shown in Figure 8a. We compare the ability of the `RGPEPlanner` to maximize reaction yield on a particular target dataset after meta-training




FIG. 7. Results of multi-fidelity optimization benchmark using the HOIP dataset from Kim *et al.*¹²³ Boxplots show the cumulative simulation cost incurred before identification of the perovskite material with the smallest HSE06 bandgap in the dataset. Each strategy is run 100 independent times. The `MultiFidelityPlanner` also has access to bandgap measurements at the GGA level (with an E_g^{GGA} to E_g^{HSE06} measurement ratio of 8:1), while the `GPPlanner` has access to E_g^{HSE06} measurements only.

on yield measurements from the other 4 datasets. For instance, if the target dataset is `buchwald_c`, the `RGPEPlanner` has access to measurements from the `buchwald_{a,b,d,e}` datasets. As a baseline, we optimize each target reaction using ATLAS' `GPPlanner`, which has no access to historical reaction data.

The following code snippet shows the problem setup. The `RGPEPlanner` or `DKTPlanner` takes an argument called `train_tasks`, which must be a list of dictionaries containing the source task data (`source_tasks` are provided for this example in the GitHub repo).

```

from olympus import Dataset, Campaign
from atlas.planners import RGPEPlanner

target_dataset = Dataset(kind='buchwald_c') # instantiate the target dataset from Olympus
source_tasks = load_my_source_tasks() # load the source task data (provided on GitHub repo)
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(target_dataset.param_space)

planner = RGPEPlanner(goal='maximize', train_tasks=source_tasks) # instantiate Atlas planner
planner.set_param_space(target_dataset.param_space)

while campaign.num_obs < 50:
    samples = planner.recommend(campaign.observations) # ask planner for batch of parameters
    for sample in samples:

```

```

measurement = target_dataset.run(sample) # measure objective of the target dataset
campaign.add_observation(sample, measurement) # tell planner about most recent observation

```

Figure 8b shows the results of this comparative experiment for 3 target reaction products: buchwald_{c,d,e}. In each case, the RGPEPlanner is able to identify higher yields with fewer objective evaluations compared to the GPlanner by using intuition gleaned from meta-training on related reaction data.

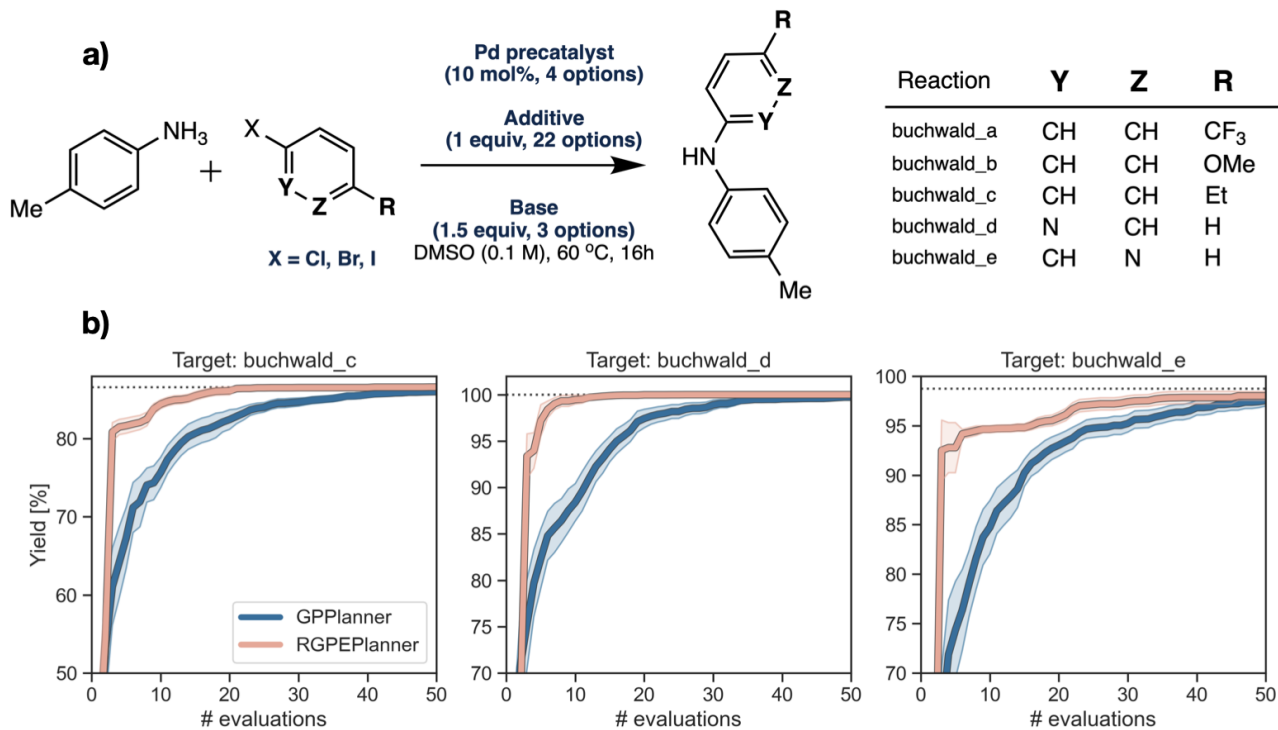


FIG. 8. a) Reaction scheme for the Buchwald-Hartwig datasets. b) Optimization traces comparing the performance of ATLAS' RGPEPlanner and GPlanner on 3 target reaction datasets. Solid traces show the mean taken over 50 independently seeded runs. Shaded regions show the 95% confidence interval. Horizontal dotted lines indicate the maximum possible yield reported by Ahneman *et al.*¹³³ for each reaction product.

H. Optimization over molecular domains

For optimization over molecular spaces, we provide a specialized GP kernel function that is compatible with all planners and is based on the Tanimoto distance kernel.^{81,92} The Tanimoto kernel is a general similarity metric^{134,135} defined for binary vectors $\mathbf{x}, \mathbf{x}' \in \{0, 1\}^d$ for $d \geq 1$ as

$$k_{\text{Tanimoto}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \frac{\langle \mathbf{x}, \mathbf{x}' \rangle}{\|\mathbf{x}\|^2 + \|\mathbf{x}'\|^2 - \langle \mathbf{x}, \mathbf{x}' \rangle}, \quad (10)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product, $\|\cdot\|$ is the Euclidean norm, and σ_f^2 is the kernel's signal variance hyperparameter. Several binary vector representations of molecules are available, but perhaps the most well known are extended-connectivity fingerprints (ECFPs).¹³⁶

ATLAS allows for use of any molecular representation based on binary vectors. Users must only specify, to the constructor of the planner, the indices that identify molecular parameters with the parameter space using the `molecular_params` argument. Molecular parameters must be of type categorical, with options corresponding to unique molecules. Users may then define their descriptors as the corresponding binary vectors. We show a simple example in which we intend to minimize the aqueous solubility of molecules in the ESOL dataset.¹³⁷

```

from olympus import Campaign, ParameterSpace, ParameterCategorical
from atlas.planners.gp.planner import GPPlanner

smiles = load_esol_smiles() # load list of smiles for molecules in ESOL (provided on GitHub repo)
ecfps = load_esol_ecfp() # load descriptors ECFP vectors for molecules in ESOL (provided on GitHub repo)

# create parameter space with one molecular parameter
param_space = ParameterSpace()
param_space.add(ParameterCategorical(name='esol', options=smiles, descriptors=ecfps))
campaign = Campaign() # define Olympus campaign object
campaign.set_param_space(param_space)

# instantiate planner with list of parameter space indices of molecular parameter(s)
planner = GPPlanner(goal='minimize', molecular_params=[0])
planner.set_param_space(param_space)

```

While it has been shown that using physicochemical descriptor-based representations of molecules in a BO setting can accelerate optimization rate,^{78,138} other studies have found that expert-crafted descriptors did not out-perform simpler representations like fingerprints or even one-hot-encodings.¹³⁹ Given the apparent dependence of optimal molecular representation on the characteristics of the optimization problem at hand, ATLAS provides users with the flexibility to represent molecular parameters in several ways.

I. Asynchronous experimental execution

In many SDL applications, researchers have access to multiple robotic or computational workers and may parallelize measurements. When performing batched BO in a setting where there is variability in measurement times for each individual experiment, it is important to operate an SDL asynchronously, where a worker starts a new experiment immediately after completion of the previous experiment.¹⁴⁰ This approach has been shown to be overall more efficient than waiting for an entire batch of experiments to complete before commencing the next batch.^{72,73}

We provide template scripts for an asynchronous SDL setup on our GitHub repo. In this example, we optimize a surface from OLYMPUS using 3 workers, each of which can perform a measurement for a single set of parameters. Workers can be assigned parameters to measure in parallel using multiprocessing, and measurement duration is set to be variable. Upon receiving a measurement, ATLAS re-trains its surrogate model, but also conditions its recommendations on pending experiments, *i.e.* those that have been assigned to a worker but whose measurement

has not completed (the reader is referred to Section III A for additional information on setting pending parameter constraints in ATLAS). This is achieved with a built-in mechanism to generate fictitious measurements, y' for pending parameters, \mathbf{x}' . Specifically, we adopt the hallucination or *kriging believer* strategy first reported by Ginsbourger *et al.*,^{72,73} which imputes the expected value of each pending parameter,

$$y' = \mathbb{E}[y|\mathbf{x}', \mathcal{D}], \quad (11)$$

where \mathcal{D} is the current dataset of observations. The fictitious measurement y' is then used to augment the dataset of observations, *i.e.* $\mathcal{D}' = \mathcal{D} \cup \{(\mathbf{x}', y')\}$. Effectively, this strategy fixes the posterior mean of the surrogate model and updates its variance in light of pending parameters. ATLAS maintains a “priority queue” of recommended parameters that is immediately updated in light of new measurements, such that parameter proposals are always informed by the most recent observations. Proposals in the priority queue are then delegated to measurement workers as they become available.

IV. EXPERIMENTAL DEMONSTRATION

In this section, we outline an experimental demonstration highlighting the use of ATLAS in a simple SDL. Specifically, we show how ATLAS is combined with CHEMOS 2.0,⁸² an SDL orchestration software, to optimize the oxidation potential of a set of metal complexes in a cyclic voltammetry (CV) experiment. Figure 9a depicts the experimental setup. The electrochemical SDL consists of two hardware parts: the automatic complexation robot module and the E-chem analyzer module, both of which are controlled remotely from CHEMOS 2.0. CHEMOS 2.0 hosts ATLAS, and iteratively sends jobs to the E-chem setup for each evaluation step of the optimization campaign. In turn, it receives the raw data and the treated oxidation potential from the instrument after execution of the CV experiment. CHEMOS 2.0 saves the raw CV data in its internal experimental database and also saves the results of the optimization campaign. Data for the experiment reported in this work has all been stored on CHEMOS 2.0.

The automatic complexation robot is a flow-based system based on a syringe pump and selection valves driven by a Python controller developed in-house. Upon receiving instructions, it runs automatic complexation by transferring designated amounts of stock solutions (metal, ligand, electrolyte, buffer and water) to the reactor, conducts a reactor mixing step, transfers the sample to the flow cell of the E-chem module, and invokes the electrochemistry measurements. A standard clean-up step is executed after the electrochemistry is finished. The E-chem analyzer consists of a flow cell equipped with a printed electrode and a low-cost potentiostat controlled by Python software. The electrochemistry measurement is invoked by CHEMOS 2.0 to measure the sample in the flow cell. In the context of this demonstration, a fixed parameter CV experiment is conducted for all samples. The raw data collected from the potentiostat is streamed to CHEMOS 2.0 for further processing, and cleaning instructions are sent to the complexation robot.

The parameter space for this experiment consists of 4 parameters, which are summarized in Table I. The ligand options are H₂O, Pyridine, and Ethylenediamine. The metal options are Silver (I) and Copper (II). The objective of the optimization is to maximize the voltage of the oxidation peak. The **GPPlanner** of ATLAS is run for 40 iterations, the first 5 of which are randomly selected initial design points. Results of the optimization experiment are shown in Figure 9b and 9c. Figure 9b shows two cyclic voltammograms collected during the optimization. The voltage peak is

selected using in-house software and is identified on the figure with a red cross and red vertical dotted line. Figure 9c shows the maximum voltage peak identified by the **GPPlanner** as a function of the number of transpired experiments.

TABLE I. Parameter space for the electrochemical SDL demonstration with ATLAS. † Discrete parameter has a stride of 1.

Parameter	Type	Range / Num Options	Description	Descriptors
Ligand	Categorical	3	Ligand identity	No
Metal	Categorical	2	Metal identity	No
# mixings	Discrete	[1 - 10] [†]	Number of pump mixing steps	N/A
Ratio	Continuous	[1.0 - 9.0]	Ligand/metal ratio	N/A

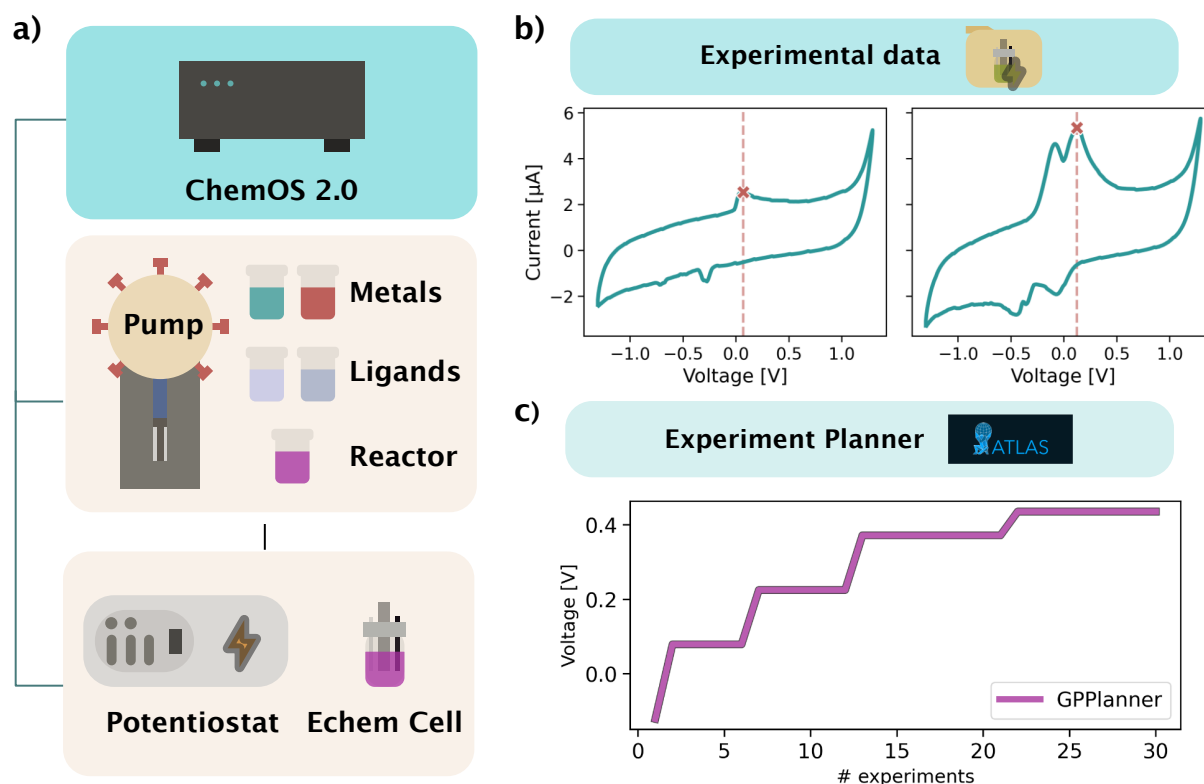


FIG. 9. a) Schematic diagram of automated electrochemical SDL setup. b) Two examples of cyclic voltammograms measured during the campaign. The oxidation peak is identified using in-house software and is marked with a red cross and vertical dotted line. c) Optimization trace showing the maximum oxidation peak voltage identified by ATLAS' **GPPlanner** as a function of the number of completed experiments.

V. CONCLUSION

In summary, we introduce ATLAS, a Bayesian optimization package with a comprehensive set of features designed to suit most experimental settings and enable model-guided optimization in SDLs. Among the capabilities currently available are optimizations over mixed-parameter, constrained, and molecular domains, in addition to supporting multi-objective and multi-fidelity optimizations, as well as robust optimization and the incorporation of past knowledge

via meta-learning. ATLAS uses Gaussian process surrogate models, and is built upon PYTORCH,⁹⁴ GPYTORCH,⁹² and BOTORCH.⁵⁸ It exposes its broad set of capabilities via the OLYMPUS⁷⁹ interface, and it integrates with CHEMOS 2.0⁸² for SDL deployment. ATLAS is an open-source software, it is distributed under the MIT permissive license, and comes with a documentation that includes examples for all case scenarios discussed in this manuscript. We expect ATLAS to be able to cover a much broader set of SDL setups and research challenges than the Bayesian optimization packages developed to date.

CODE AND DATA AVAILABILITY

ATLAS is available on GitHub at <https://github.com/aspuru-guzik-group/atlas> under an MIT license. The measurements generated in the electrochemical SDL demonstration have been added to the OLYMPUS package as an emulated Dataset called `electrochem`, on which users may benchmark experiment planning strategies (https://github.com/aspuru-guzik-group/olympus/tree/dev/src/olympus/datasets/dataset_electrochem) CHEMOS 2.0 is available on GitHub at <https://github.com/malcolmsimgithub/ChemOS2.0> under an MIT license.

ACKNOWLEDGMENTS

The authors thank Dr. Felix Strieth-Kalthoff, Dr. Martin Seifrid, Dr. Shengyang Sun, Dr. Roger Grosse, Dr. Robert Black, Guodong Zhang, Gary Tom, and Erfan Fathei for contribution to valuable discussion. R.J.H. gratefully acknowledges the Natural Sciences and Engineering Research Council of Canada (NSERC) for provision of the Post-graduate Scholarships-Doctoral Program (PGSD3-534584-2019), as well as support from the Vector Institute. C.A. acknowledges an NSERC Discovery grant (RGPIN-2022-04910). A.A.-G. acknowledges support from the Canada 150 Research Chairs program and CIFAR, as well as the generous support of Dr. Anders G. Frøseth. This research was undertaken thanks in part to funding provided to the University of Toronto's Acceleration Consortium from the Canada First Research Excellence Fund (CFREF). Computations reported in this work were performed on the computing clusters of the Vector Institute and on the Niagara supercomputer at the SciNet HPC Consortium.^{141,142} Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute. SciNet is funded by the Canada Foundation for Innovation, the Government of Ontario, Ontario Research Fund - Research Excellence, and by the University of Toronto.

CONFLICTS OF INTEREST

A.A.-G is a founding member of Kebotix, Inc. R.J.H, Z.B., P.B., C.A. and A.A.-G. are founding members of a new company, 15073383 Canada Inc.

- [1] F. Häse, L. M. Roch, and A. Aspuru-Guzik, "Next-Generation Experimentation with Self-Driving Laboratories," *Trends in Chemistry*, vol. 1, pp. 282–291, June 2019.
- [2] E. Stach, B. DeCost, A. G. Kusne, J. Hattrick-Simpers, K. A. Brown, K. G. Reyes, J. Schrier, S. Billinge, T. Buonassisi, I. Foster, C. P. Gomes, J. M. Gregoire, A. Mehta, J. Montoya, E. Olivetti, C. Park, E. Rotenberg, S. K. Saikin, S. Smullin, V. Stanev, and B. Maruyama, "Autonomous experimentation systems for materials development: A community perspective," *Matter*, vol. 4, pp. 2702–2726, Sept. 2021.
- [3] C. W. Coley, N. S. Eyke, and K. F. Jensen, "Autonomous Discovery in the Chemical Sciences Part I: Progress," *Angewandte Chemie International Edition*, vol. 59, no. 51, pp. 22858–22893, 2020.
- [4] C. W. Coley, N. S. Eyke, and K. F. Jensen, "Autonomous Discovery in the Chemical Sciences Part II: Outlook," *Angewandte Chemie International Edition*, vol. 59, no. 52, pp. 23414–23436, 2020.
- [5] M. M. Flores-Leonar, L. M. Mejía-Mendoza, A. Aguilar-Granda, B. Sanchez-Lengeling, H. Tribukait, C. Amador-Bedolla, and A. Aspuru-Guzik, "Materials Acceleration Platforms: On the way to autonomous experimentation," *Current Opinion in Green and Sustainable Chemistry*, vol. 25, p. 100370, 2020.
- [6] H. S. Stein and J. M. Gregoire, "Progress and prospects for accelerating materials science with automated and autonomous workflows," *Chemical Science*, vol. 10, no. 42, pp. 9640–9649, 2019.
- [7] J. Yano, K. J. Gaffney, J. Gregoire, L. Hung, A. Ourmazd, J. Schrier, J. A. Sethian, and F. M. Toma, "The case for data science in experimental chemistry: examples and recommendations," *Nat Rev Chem*, vol. 6, pp. 357–370, Apr. 2022.
- [8] J. P. McMullen and K. F. Jensen, "An automated microfluidic system for online optimization in chemical synthesis," *Organic process research & development*, vol. 14, no. 5, pp. 1169–1176, 2010.
- [9] D. E. Fitzpatrick, C. Battilocchio, and S. V. Ley, "A novel internet-based reaction monitoring, control and autonomous self-optimization platform for chemical synthesis," *Organic Process Research & Development*, vol. 20, no. 2, pp. 386–394, 2016.
- [10] D. Cortés-Borda, E. Wimmer, B. Gouilleux, E. Barré, N. Oger, L. Goulamaly, L. Peault, B. Charrier, C. Truchet, P. Giraudeau, *et al.*, "An autonomous self-optimizing flow reactor for the synthesis of natural product carpanone," *The Journal of organic chemistry*, vol. 83, no. 23, pp. 14286–14299, 2018.
- [11] B. E. Walker, J. H. Bannock, A. M. Nightingale, and J. C. deMello, "Tuning reaction products by constrained optimisation," *Reaction Chemistry & Engineering*, vol. 2, no. 5, pp. 785–798, 2017.
- [12] S. Krishnadasan, R. Brown, A. Demello, and J. Demello, "Intelligent routes to the controlled synthesis of nanoparticles," *Lab on a Chip*, vol. 7, no. 11, pp. 1434–1441, 2007.
- [13] L. M. Baumgartner, C. W. Coley, B. J. Reizman, K. W. Gao, and K. F. Jensen, "Optimum catalyst selection over continuous and discrete process variables with a single droplet microfluidic reaction platform," *Reaction Chemistry & Engineering*, vol. 3, no. 3, pp. 301–311, 2018.
- [14] A. Schweidtmann, A. Clayton, N. Holmes, E. Bradford, R. Bourne, and A. Lapkin, "Machine learning meets continuous flow chemistry: Automated optimization towards the pareto front of multiple objectives," *Chem. Eng. J.*, pp. 177–282, 2018.

- [15] A.-C. Bédard, A. Adamo, K. C. Aroh, M. G. Russell, A. A. Bedermann, J. Torosian, B. Yue, K. F. Jensen, and T. F. Jamison, “Reconfigurable system for automated optimization of diverse chemical reactions,” *Science*, vol. 361, pp. 1220–1225, Sept. 2018.
- [16] M. Christensen, L. P. E. Yunker, F. Adedeji, F. Häse, L. M. Roch, T. Gensch, G. dos Passos Gomes, T. Zepel, M. S. Sigman, A. Aspuru-Guzik, and J. E. Hein, “Data-science driven autonomous process optimization,” *Communications Chemistry*, vol. 4, pp. 1–12, Aug. 2021.
- [17] A. M. K. Nambiar, C. P. Breen, T. Hart, T. Kulesza, T. F. Jamison, and K. F. Jensen, “Bayesian Optimization of Computer-Proposed Multistep Synthetic Routes on an Automated Robotic Flow Platform,” *ACS Central Science*, vol. 8, pp. 825–836, June 2022.
- [18] J. M. Granda, L. Donina, V. Dragone, D.-L. Long, and L. Cronin, “Controlling an organic synthesis robot with machine learning to search for new reactivity,” *Nature*, vol. 559, no. 7714, pp. 377–381, 2018.
- [19] P. Nikolaev, D. Hooper, F. Webber, R. Rao, K. Decker, M. Krein, J. Poleski, R. Barto, and B. Maruyama, “Autonomy in materials research: a case study in carbon nanotube growth,” *npj Computational Materials*, vol. 2, no. 1, pp. 1–6, 2016.
- [20] K. Vaddi, H. T. Chiang, and L. D. Pozzo, “Autonomous retrosynthesis of gold nanoparticles via spectral shape matching,” *Digital Discovery*, vol. 1, pp. 502–510, Aug. 2022. Publisher: RSC.
- [21] R. J. Hickman, P. Bannigan, Z. Bao, A. Aspuru-Guzik, and C. Allen, “Self-driving laboratories: A paradigm shift in nanomedicine development,” *Matter*, Mar. 2023.
- [22] A. Deshwal, C. M. Simon, and J. Rao Doppa, “Bayesian optimization of nanoporous materials,” *Molecular Systems Design & Engineering*, vol. 6, no. 12, pp. 1066–1086, 2021.
- [23] B. P. MacLeod, F. G. L. Parlane, T. D. Morrissey, F. Häse, L. M. Roch, K. E. Dettelbach, R. Moreira, L. P. E. Yunker, M. B. Rooney, J. R. Deeth, V. Lai, G. J. Ng, H. Situ, R. H. Zhang, M. S. Elliott, T. H. Haley, D. J. Dvorak, A. Aspuru-Guzik, J. E. Hein, and C. P. Berlinguette, “Self-driving laboratory for accelerated discovery of thin-film materials,” *Science Advances*, vol. 6, p. eaaz8867, May 2020.
- [24] B. P. MacLeod, F. G. L. Parlane, C. C. Rupnow, K. E. Dettelbach, M. S. Elliott, T. D. Morrissey, T. H. Haley, O. Proskurin, M. B. Rooney, N. Taherimaksousi, D. J. Dvorak, H. N. Chiu, C. E. B. Waizenegger, K. Ocean, M. Mokhtari, and C. P. Berlinguette, “A self-driving laboratory advances the Pareto front for material properties,” *Nature Communications*, vol. 13, p. 995, Feb. 2022.
- [25] N. T. P. Hartono, M. Ani Najeeb, Z. Li, P. W. Nega, C. A. Fleming, X. Sun, E. M. Chan, A. Abate, A. J. Norquist, J. Schrier, and T. Buonassisi, “Principled Exploration of Bipyridine and Terpyridine Additives to Promote Methylammonium Lead Iodide Perovskite Crystallization,” *Crystal Growth & Design*, vol. 22, pp. 5424–5431, Sept. 2022. Publisher: American Chemical Society.
- [26] S. Sun, A. Tiihonen, F. Oviedo, Z. Liu, J. Thapa, Y. Zhao, N. T. P. Hartono, A. Goyal, T. Heumueller, C. Batali, *et al.*, “A data fusion approach to optimize compositional stability of halide perovskites,” *Matter*, vol. 4, no. 4, pp. 1305–1322, 2021.
- [27] B. Burger, P. M. Maffettone, V. V. Gusev, C. M. Aitchison, Y. Bai, X. Wang, X. Li, B. M. Alston, B. Li, R. Clowes, N. Rankin, B. Harris, R. S. Sprick, and A. I. Cooper, “A mobile robotic chemist,” *Nature*, vol. 583, pp. 237–241, July 2020.
- [28] M. J. Tamasi, R. A. Patel, C. H. Borca, S. Kosuri, H. Mugnier, R. Upadhy, N. S. Murthy, M. A. Webb, and A. J. Gormley, “Machine Learning on a Robotic Platform for the Design of Polymer–Protein Hybrids,” *Advanced Materials*, vol. 34, no. 30, p. 2201809, 2022.
- [29] M. J. Tamasi and A. J. Gormley, “Biologic formulation in a self-driving biomaterials lab,” *Cell Reports Physical Science*, vol. 3, p. 101041, Sept. 2022.

- [30] M. M. Noack, K. G. Yager, M. Fukuto, G. S. Doerk, R. Li, and J. A. Sethian, “A kriging-based approach to autonomous experimentation with applications to X-ray scattering,” *Scientific reports*, vol. 9, no. 1, pp. 1–19, 2019.
- [31] B. Rohr, H. S. Stein, D. Guevarra, Y. Wang, J. A. Haber, M. Aykol, S. K. Suram, and J. M. Gregoire, “Benchmarking the acceleration of materials discovery by sequential learning,” *Chemical Science*, vol. 11, no. 10, pp. 2696–2706, 2020.
- [32] R. A. Fisher, *The design of experiments*. Oliver and Boyd; Edinburgh; London, 1937.
- [33] G. E. P. Box, J. S. Hunter, and W. G. Hunter, *Statistics for experimenters: design, innovation and discovery*, vol. 2. 2005.
- [34] M. J. Anderson and P. J. Whitcomb, *DOE simplified: practical tools for effective experimentation*. CRC Press, 2016.
- [35] A. Lucia and J. Xu, “Chemical process optimization using newton-like methods,” *Computers & chemical engineering*, vol. 14, no. 2, pp. 119–138, 1990.
- [36] I. Rechenberg, “Evolutionsstrategien,” in *Simulationenmethoden in der Medizin und Biologie*, pp. 83–114, Springer, 1978.
- [37] H.-P. Schwefel, “Evolutionsstrategien für die numerische optimierung,” in *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie*, pp. 123–176, Springer, 1977.
- [38] G. Zames, N. Ajlouni, N. Ajlouni, N. Ajlouni, J. Holland, W. Hills, and D. Goldberg, “Genetic algorithms in search, optimization and machine learning,” *Information Technology Journal*, vol. 3, no. 1, pp. 301–302, 1981.
- [39] J. R. Koza and J. R. Koza, *Genetic programming: on the programming of computers by means of natural selection*, vol. 1. MIT press, 1992.
- [40] M. Srinivas and L. M. Patnaik, “Genetic algorithms: A survey,” *computer*, vol. 27, no. 6, pp. 17–26, 1994.
- [41] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, “Natural evolution strategies,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 949–980, 2014.
- [42] Z. Zhou, X. Li, and R. N. Zare, “Optimizing Chemical Reactions with Deep Reinforcement Learning,” *ACS Central Science*, vol. 3, pp. 1337–1344, Dec. 2017.
- [43] C. Beeler, S. G. Subramanian, K. Sprague, N. Chatti, C. Bellinger, M. Shahen, N. Paquin, M. Baula, A. Dawit, Z. Yang, X. Li, M. Crowley, and I. Tamblyn, “Chemgymrl: An interactive framework for reinforcement learning for digital chemistry,” 2023.
- [44] J. Močkus, “On bayesian methods for seeking the extremum,” in *Optimization techniques IFIP technical conference*, pp. 400–404, Springer, 1975.
- [45] J. Mockus, V. Tiesis, and A. Zilinskas, “The application of bayesian methods for seeking the extremum,” *Towards global optimization*, vol. 2, no. 117-129, p. 2, 1978.
- [46] J. Mockus, *Bayesian approach to global optimization: theory and applications*, vol. 37. Springer Science & Business Media, 2012.
- [47] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [48] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and G. Varoquaux, “API design for machine learning software: experiences from the scikit-learn project,” in *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pp. 108–122, 2013.
- [49] The GPyOpt authors, “GPyOpt: A bayesian optimization framework in python.” <http://github.com/SheffieldML/GPyOpt>, 2016.
- [50] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, “Algorithms for hyper-parameter optimization,” in *Advances in neural information processing systems*, pp. 2546–2554, 2011.

- [51] J. S. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization,” *Journal of Machine Learning Research*, vol. 13, no. Feb, pp. 281–305, 2012.
- [52] J. S. Bergstra, D. Yamins, and D. D. Cox, “Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures,” *JMLR*, 2013.
- [53] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins, and D. D. Cox, “Hyperopt: a python library for model selection and hyperparameter optimization,” *Computational Science & Discovery*, vol. 8, no. 1, p. 014008, 2015.
- [54] M. Lindauer, K. Eggenberger, M. Feurer, A. Biedenkapp, D. Deng, C. Benjamins, T. Ruhkopf, R. Sass, and F. Hutter, “Smac3: A versatile bayesian optimization package for hyperparameter optimization,” 2021.
- [55] M. Lindauer, K. Eggenberger, M. Feurer, S. Falkner, A. Biedenkapp, and F. Hutter, “Smac v3: Algorithm configuration in python.” <https://github.com/automl/SMAC3>, 2017.
- [56] K. Kandasamy, K. R. Vysyaraju, W. Neiswanger, B. Paria, C. R. Collins, J. Schneider, B. Póczos, and E. P. Xing, “Tuning hyperparameters without grad students: Scalable and robust bayesian optimisation with dragonfly,” *J. Mach. Learn. Res.*, vol. 21, no. 81, pp. 1–27, 2020.
- [57] A. I. Cowen-Rivers, W. Lyu, R. Tutunov, Z. Wang, A. Grosnit, R. R. Griffiths, H. Jianye, J. Wang, and H. B. Ammar, “An empirical study of assumptions in bayesian optimisation,” *arXiv preprint arXiv:2012.03826*, 2020.
- [58] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, “BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization,” in *Advances in Neural Information Processing Systems 33*, 2020.
- [59] The Ax authors, “Ax: Adaptive experimentation platform.” <https://github.com/facebook/Ax>, 2023.
- [60] X. Song, S. Perel, C. Lee, G. Kochanski, and D. Golovin, “Open source vizier: Distributed infrastructure and api for reliable and flexible black-box optimization,” in *Automated Machine Learning Conference, Systems Track (AutoML-Conf Systems)*, 2022.
- [61] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. Karro, and D. Sculley, “Google vizier: A service for black-box optimization,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*, pp. 1487–1495, ACM, 2017.
- [62] M. Aldeghi, F. Häse, R. J. Hickman, I. Tamblin, and A. Aspuru-Guzik, “Golem: an algorithm for robust experiment and process optimization,” *Chemical Science*, Oct. 2021. Publisher: The Royal Society of Chemistry.
- [63] S. Daulton, S. Cakmak, M. Balandat, M. A. Osborne, E. Zhou, and E. Bakshy, “Robust multi-objective Bayesian optimization under input noise,” in *Proceedings of the 39th International Conference on Machine Learning* (K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, eds.), vol. 162 of *Proceedings of Machine Learning Research*, pp. 4831–4866, PMLR, 17–23 Jul 2022.
- [64] R. J. Hickman, M. Aldeghi, F. Häse, and A. Aspuru-Guzik, “Bayesian optimization with known experimental and design constraints for chemistry applications,” *Digital Discovery*, Sept. 2022. Publisher: RSC.
- [65] S. G. Baird, J. R. Hall, and T. D. Sparks, “Compactness matters: Improving Bayesian optimization efficiency of materials formulations through invariant search spaces,” *Computational Materials Science*, vol. 224, p. 112134, May 2023.
- [66] J. Snoek, H. Larochelle, and R. P. Adams, “Practical Bayesian Optimization of Machine Learning Algorithms,” *arXiv:1206.2944 [cs, stat]*, Aug. 2012. arXiv: 1206.2944.
- [67] M. A. Gelbart, J. Snoek, and R. P. Adams, “Bayesian Optimization with Unknown Constraints,” *arXiv:1403.5607 [cs, stat]*, Mar. 2014. arXiv: 1403.5607.
- [68] R. B. Gramacy and H. K. H. Lee, “Optimization Under Unknown Constraints,” *arXiv:1004.4027 [stat]*, July 2010. arXiv: 1004.4027.
- [69] C. Antonio, “Sequential model based optimization of partially defined functions under unknown constraints,” *Journal of Global Optimization*, vol. 79, pp. 281–303, Feb. 2021.

- [70] Y. K. Wakabayashi, T. Otsuka, Y. Krockenberger, H. Sawada, Y. Taniyasu, and H. Yamamoto, “Bayesian optimization with experimental failure for high-throughput materials growth,” *npj Computational Materials*, vol. 8, no. 180, 2022.
- [71] “Bayesian optimization with active learning of design constraints using an entropy-based approach,” *npj Computational Materials*, vol. 9, p. 49, Apr. 2023.
- [72] D. Ginsbourger, J. Janusevskis, and R. L. Riche, “Dealing with asynchronicity in parallel Gaussian Process based global optimization,” 2011.
- [73] T. Desautels, A. Krause, and J. W. Burdick, “Parallelizing Exploration-Exploitation Tradeoffs in Gaussian Process Bandit Optimization,” *Journal of Machine Learning Research*, vol. 15, no. 119, pp. 4053–4103, 2014.
- [74] B. J. Shields, J. Stevens, J. Li, M. Parasram, F. Damani, J. I. M. Alvarado, J. M. Janey, R. P. Adams, and A. G. Doyle, “Bayesian reaction optimization as a tool for chemical synthesis,” *Nature*, vol. 590, pp. 89–96, Feb. 2021.
- [75] J. G. Torres, S. H. Lau, P. Anchuri, J. Stevens, J. Tabora, J. Li, A. Borovika, R. Adams, and A. Doyle, “A Multi-Objective Active Learning Platform and Web App for Reaction Optimization,” Aug. 2022.
- [76] K. C. Felton, J. G. Rittig, and A. A. Lapkin, “Summit: Benchmarking Machine Learning Methods for Reaction Optimization,” *Chemistry-Methods*, vol. 1, no. 2, pp. 116–122, 2021.
- [77] F. Häse, L. M. Roch, C. Kreisbeck, and A. Aspuru-Guzik, “Phoenix: A Bayesian Optimizer for Chemistry,” *ACS Central Science*, vol. 4, pp. 1134–1145, Sept. 2018.
- [78] F. Häse, M. Aldeghi, R. J. Hickman, L. M. Roch, and A. Aspuru-Guzik, “Gryffin: An algorithm for Bayesian optimization of categorical variables informed by expert knowledge,” *Applied Physics Reviews*, vol. 8, p. 031406, Sept. 2021.
- [79] F. Häse, M. Aldeghi, R. J. Hickman, L. M. Roch, M. Christensen, E. Liles, J. E. Hein, and A. Aspuru-Guzik, “Olympus: a benchmarking framework for noisy optimization and experiment planning,” *Machine Learning: Science and Technology*, vol. 2, p. 035021, July 2021. Publisher: IOP Publishing.
- [80] R. Hickman, P. Parakh, A. Cheng, Q. Ai, J. Schrier, M. Aldeghi, and A. Aspuru-Guzik, “Olympus, enhanced: benchmarking mixed-parameter and multi-objective optimization in chemistry and materials science,” May 2023.
- [81] R.-R. Griffiths, L. Klarner, H. B. Moss, A. Ravuri, S. Truong, B. Rankovic, Y. Du, A. Jamasb, J. Schwartz, A. Tripp, G. Kell, A. Bourached, A. Chan, J. Moss, C. Guo, A. A. Lee, P. Schwaller, and J. Tang, “Gauche: A library for gaussian processes in chemistry,” 2022.
- [82] M. Sim, M. G. Vakili, F. Strieth-Kalthoff, H. Hao, R. Hickman, S. Miret, S. Pablo-García, and A. Aspuru-Guzik, “ChemOS 2.0: an orchestration architecture for chemical self-driving laboratories,” Aug. 2023.
- [83] D. Kraft *et al.*, “A software package for sequential quadratic programming,” 1988.
- [84] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv:1412.6980 [cs]*, Jan. 2017. arXiv: 1412.6980.
- [85] J. Blank and K. Deb, “pymoo: Multi-objective optimization in python,” *IEEE Access*, vol. 8, pp. 89497–89509, 2020.
- [86] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné, “DEAP: Evolutionary algorithms made easy,” *Journal of Machine Learning Research*, vol. 13, pp. 2171–2175, 2012.
- [87] F.-M. De Rainville, F.-A. Fortin, M.-A. Gardner, M. Parizeau, and C. Gagné, “Deap: A python framework for evolutionary algorithms,” in *Proceedings of the 14th annual conference companion on Genetic and evolutionary computation*, pp. 85–92, 2012.
- [88] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. Adaptive computation and machine learning, MIT Press, 2006.
- [89] R. M. Neal, *Bayesian Learning for Neural Networks*, vol. 118 of *Lecture Notes in Statistics*. New York, NY: Springer, 1996.
- [90] C. Rasmussen and Z. Ghahramani, “Occam’s Razor,” in *Advances in Neural Information Processing Systems*, vol. 13, MIT Press, 2000.

- [91] J. Hensman, A. G. de G. Matthews, and Z. Ghahramani, “Scalable variational gaussian process classification,” in *International Conference on Artificial Intelligence and Statistics*, 2014.
- [92] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, “Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration,” in *Advances in Neural Information Processing Systems*, 2018.
- [93] G. P. Szego and L. C. W. Dixon, eds., *Towards global optimisation 2*. Amsterdam ; New York : New York: North-Holland Pub. Co. ; sole distributors for the U.S.A. and Canada, Elsevier North-Holland, 1978.
- [94] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, Curran Associates, Inc., 2019.
- [95] The Atlas authors, “Atlas: A brain for self-driving laboratories.” <https://github.com/aspuru-guzik-group/atlas>, 2023.
- [96] P. Vellanki, S. Rana, S. Gupta, D. Rubin, A. Sutti, T. Dorin, M. Height, P. Sanders, and S. Venkatesh, “Process-constrained batch Bayesian optimisation,” in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.
- [97] E. Soedarmadji, H. S. Stein, S. K. Suram, D. Guevarra, and J. M. Gregoire, “Tracking materials science data lineage to manage millions of materials experiments and analyses,” *npj Computational Materials*, vol. 5, pp. 1–9, July 2019.
- [98] H. S. Stein, D. Guevarra, A. Shinde, R. J. R. Jones, J. M. Gregoire, and J. A. Haber, “Functional mapping reveals mechanistic clusters for OER catalysis across (Cu–Mn–Ta–Co–Sn–Fe)Ox composition and pH space,” *Materials Horizons*, vol. 6, pp. 1251–1258, July 2019.
- [99] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural networks,” 2015.
- [100] F. Häse, L. M. Roch, and A. Aspuru-Guzik, “Chimera: enabling hierarchy based multi-objective optimization for self-driving laboratories,” *Chemical Science*, vol. 9, no. 39, pp. 7642–7655, 2018.
- [101] E. Zitzler and L. Thiele, “Multiobjective optimization using evolutionary algorithms — A comparative case study,” in *Parallel Problem Solving from Nature — PPSN V* (A. E. Eiben, T. Bäck, M. Schoenauer, and H.-P. Schwefel, eds.), Lecture Notes in Computer Science, (Berlin, Heidelberg), pp. 292–301, Springer, 1998.
- [102] J. Knowles, D. Corne, and M. Fleischer, “Bounded archiving using the lebesgue measure,” in *The 2003 Congress on Evolutionary Computation, 2003. CEC '03.*, vol. 4, pp. 2490–2497 Vol.4, 2003.
- [103] M. Li and X. Yao, “Quality evaluation of solution sets in multiobjective optimisation: A survey,” *ACM Comput. Surv.*, vol. 52, mar 2019.
- [104] A. P. Guerreiro, C. M. Fonseca, and L. Paquete, “The Hypervolume Indicator: Problems and Algorithms,” *ACM Computing Surveys*, vol. 54, pp. 1–42, July 2021. arXiv:2005.00515 [cs].
- [105] J. Knowles and E. J. Hughes, “Multiobjective optimization on a budget of 250 evaluations,” in *Evolutionary Multi-Criterion Optimization (EMO-2005)* (C. Coello et al, ed.), vol. 3410 of LNCS, Springer-Verlag, 2005.
- [106] J. Knowles, “Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems,” *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 50–66, 2006.
- [107] I. Y. Kim and O. L. de Weck, “Adaptive weighted sum method for multiobjective optimization: a new method for Pareto front generation,” *Structural and Multidisciplinary Optimization*, vol. 31, pp. 105–116, Feb. 2006.
- [108] C. A. C. Coello, S. Gonzalez, L. Brambila, Josue, F. Gamboa, M. G. C. Tapia, R. Hernandez, N. Gomez, and mez, “Evolutionary multiobjective optimization: open research areas and some challenges lying ahead.,” *Complex & Intelligent Systems*, vol. 6, pp. 221–237, July 2020. Publisher: Springer.
- [109] M. Seifrid, R. J. Hickman, A. Aguilar-Granda, C. Lavigne, J. Vestfrid, T. C. Wu, T. Gaudin, E. J. Hopkins, and A. Aspuru-Guzik, “Routescore: Punching the Ticket to More Efficient Materials Development,” July 2021.

- [110] B. Settles, "Active Learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 6, pp. 1–114, June 2012. Publisher: Morgan & Claypool Publishers.
- [111] N. H. Angello, V. Rathore, W. Beker, A. Wolos, E. R. Jira, R. Roszak, T. C. Wu, C. M. Schroeder, A. Aspuru-Guzik, B. A. Grzybowski, and M. D. Burke, "Closed-loop optimization of general reaction conditions for heteroaryl Suzuki-Miyaura coupling," *Science*, vol. 378, pp. 399–405, Oct. 2022. Publisher: American Association for the Advancement of Science.
- [112] B. J. Reizman, Y.-M. Wang, S. L. Buchwald, and K. F. Jensen, "Suzuki-miyaura cross-coupling optimization enabled by automated feedback," *Reaction chemistry & engineering*, vol. 1, no. 6, pp. 658–666, 2016.
- [113] C. Fare, P. Fenner, M. Benatan, A. Varsi, and E. O. Pyzer-Knapp, "A multi-fidelity machine learning approach to high throughput materials screening," *npj Computational Materials*, vol. 8, pp. 1–9, Dec. 2022.
- [114] D. Bash, Y. Cai, V. Chellappan, S. L. Wong, X. Yang, P. Kumar, J. D. Tan, A. Abutaha, J. J. Cheng, Y.-F. Lim, S. I. P. Tian, Z. Ren, F. Mekki-Berrada, W. K. Wong, J. Xie, J. Kumar, S. A. Khan, Q. Li, T. Buonassisi, and K. Hippalgaonkar, "Multi-Fidelity High-Throughput Optimization of Electrical Conductivity in P3HT-CNT Composites," *Advanced Functional Materials*, vol. n/a, no. n/a, p. 2102606.
- [115] A. Patra, R. Batra, A. Chandrasekaran, C. Kim, T. D. Huan, and R. Ramprasad, "A multi-fidelity information-fusion approach to machine learn and predict polymer bandgap," *Computational Materials Science*, vol. 172, p. 109286, Feb. 2020.
- [116] A. Tran, J. Tranchida, T. Wildey, and A. P. Thompson, "Multi-fidelity machine-learning with uncertainty quantification and Bayesian optimization for materials design: Application to ternary random alloys," *The Journal of Chemical Physics*, vol. 153, p. 074705, Aug. 2020.
- [117] A. Tran, T. Wildey, and S. McCann, "sMF-BO-2CoGP: A Sequential Multi-Fidelity Constrained Bayesian Optimization Framework for Design Applications," *Journal of Computing and Information Science in Engineering*, vol. 20, June 2020.
- [118] N. Gantzler, A. Deshwal, J. R. Doppa, and C. Simon, "Multi-fidelity Bayesian Optimization of Covalent Organic Frameworks for Xenon/Krypton Separations," June 2023.
- [119] A. E. Gongora, K. L. Snapp, E. Whiting, P. Riley, K. G. Reyes, E. F. Morgan, and K. A. Brown, "Using simulation to accelerate autonomous experimentation: A case study using mechanics," *iScience*, vol. 24, p. 102262, Apr. 2021.
- [120] M. Poloczek, J. Wang, and P. Frazier, "Multi-Information Source Optimization," in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.
- [121] J. Wu, S. Toscano-Palmerin, P. I. Frazier, and A. G. Wilson, "Practical Multi-fidelity Bayesian Optimization for Hyperparameter Tuning," Mar. 2019. arXiv:1903.04703 [cs, math, stat].
- [122] D. Huang, T. T. Allen, W. I. Notz, and R. A. Miller, "Sequential kriging optimization using multiple-fidelity evaluations," *Structural and Multidisciplinary Optimization*, vol. 32, pp. 369–382, Nov. 2006.
- [123] C. Kim, T. Doan Huan, S. Krishnan, and R. Ramprasad, "A hybrid organic-inorganic perovskite dataset," *Scientific Data*, vol. 4, no. 170057, pp. 1–11, 2017.
- [124] J. P. Perdew, "Density functional theory and the band gap problem," *International Journal of Quantum Chemistry*, vol. 28, no. S19, pp. 497–523, 1985.
- [125] J. Heyd, G. E. Scuseria, and M. Ernzerhof, "Hybrid functionals based on a screened Coulomb potential," *The Journal of Chemical Physics*, vol. 118, pp. 8207–8215, Apr. 2003.
- [126] A. V. Krukau, O. A. Vydrov, A. F. Izmaylov, and G. E. Scuseria, "Influence of the exchange screening parameter on the performance of screened hybrid functionals," *The Journal of Chemical Physics*, vol. 125, p. 224106, Dec. 2006.
- [127] M. Feurer, B. Letham, and E. Bakshy, "Scalable Meta-Learning for Bayesian Optimization using Ranking-Weighted Gaussian Process Ensembles." 2018.
- [128] M. Feurer, B. Letham, F. Hutter, and E. Bakshy, "Practical Transfer Learning for Bayesian Optimization," arXiv:1802.02219 [cs, stat], Apr. 2021. arXiv: 1802.02219.

- [129] M. Patacchiola, J. Turner, E. J. Crowley, M. O' Boyle, and A. J. Storkey, "Bayesian Meta-Learning for the Few-Shot Setting via Deep Kernels," in *Advances in Neural Information Processing Systems*, vol. 33, pp. 16108–16118, Curran Associates, Inc., 2020.
- [130] M. Wistuba, N. Schilling, and L. Schmidt-Thieme, "Scalable Gaussian process-based transfer surrogates for hyperparameter optimization," *Machine Learning*, vol. 107, pp. 43–78, Jan. 2018.
- [131] C. J. Taylor, K. C. Felton, D. Wigh, M. I. Jeraal, R. Grainger, G. Chessari, C. N. Johnson, and A. A. Lapkin, "Accelerated Chemical Reaction Optimization Using Multi-Task Learning," *ACS Central Science*, vol. 9, pp. 957–968, May 2023.
- [132] R. J. Hickman, J. Ruža, H. Tribukait, L. M. Roch, and A. García-Durán, "Equipping data-driven experiment planning for Self-driving Laboratories with semantic memory: case studies of transfer learning in chemical reaction optimization," *Reaction Chemistry & Engineering*, June 2023.
- [133] D. T. Ahneman, J. G. Estrada, S. Lin, S. D. Dreher, and A. G. Doyle, "Predicting reaction performance in C–N cross-coupling using machine learning," *Science*, vol. 360, pp. 186–190, Apr. 2018.
- [134] J. C. Gower, "A General Coefficient of Similarity and Some of Its Properties," *Biometrics*, vol. 27, no. 4, pp. 857–871, 1971.
- [135] L. Ralaivola, S. J. Swamidass, H. Saigo, and P. Baldi, "Graph kernels for chemical informatics," *Neural Networks: The Official Journal of the International Neural Network Society*, vol. 18, pp. 1093–1110, Oct. 2005.
- [136] D. Rogers and M. Hahn, "Extended-Connectivity Fingerprints," *Journal of Chemical Information and Modeling*, vol. 50, pp. 742–754, May 2010.
- [137] J. S. Delaney, "ESOL: Estimating Aqueous Solubility Directly from Molecular Structure," *Journal of Chemical Information and Computer Sciences*, vol. 44, pp. 1000–1005, May 2004.
- [138] G. Tom, R. J. Hickman, A. Zinzuwadia, A. Mohajeri, B. Sanchez-Lengeling, and A. Aspuru-Guzik, "Calibration and generalizability of probabilistic models on low-data chemical datasets with DIONYSUS," *Digital Discovery*, vol. 2, no. 3, pp. 759–774, 2023.
- [139] A. Pomberger, A. A. Pedrina McCarthy, A. Khan, S. Sung, C. J. Taylor, M. J. Gaunt, L. Colwell, D. Walz, and A. A. Lapkin, "The effect of chemical representation on active machine learning towards closed-loop optimization," *React. Chem. Eng.*, pp. –, 2022.
- [140] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos, "Parallelised Bayesian Optimisation via Thompson Sampling," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pp. 133–142, PMLR, Mar. 2018. ISSN: 2640-3498.
- [141] M. Ponce, R. van Zon, S. Northrup, D. Gruner, J. Chen, F. Ertinaz, A. Fedoseev, L. Groer, F. Mao, B. C. Mundim, *et al.*, "Deploying a top-100 supercomputer for large parallel workloads: The niagara supercomputer," in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, pp. 1–8, 2019.
- [142] C. Loken, D. Gruner, L. Groer, R. Peltier, N. Bunn, M. Craig, T. Henriques, J. Dempsey, C.-H. Yu, J. Chen, *et al.*, "Scinet: lessons learned from building a power-efficient top-20 system and data centre," in *Journal of Physics-Conference Series*, vol. 256, p. 012026, 2010.