

# Bayesian Optimization for Accelerated Exploration of Chemical Space


**José Miguel Hernández-Lobato**

Department of Engineering  
University of Cambridge

<http://jmh1.org>, [jmh233@cam.ac.uk](mailto:jmh233@cam.ac.uk)

## Drug and material design

**Goal:** find novel molecules that optimally fulfill various metrics.




About  $10^8$  compounds in databases, potential ones:  $10^{20} - 10^{60}$ .

### Challenges:

- Evaluating molecular properties is slow and expensive.
- Chemical space is huge.

## Drug and material design

**Goal:** find novel molecules that optimally fulfill various metrics.



About  $10^8$  compounds in databases, potential ones:  $10^{20} - 10^{60}$ .

### Challenges:


- Evaluating molecular properties is slow and expensive.
- Chemical space is huge.

**Bayesian optimization** can accelerate the search.

**Bayesian optimization** aims to efficiently optimize **black-box** functions:

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

**No gradients**, observations may be **corrupted by noise**.




Black-box queries are **very expensive** (time, economic cost, etc...).



**Bayesian optimization** aims to efficiently optimize **black-box** functions:


$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

**No gradients**, observations may be **corrupted by noise**.



Black-box queries are **very expensive** (time, economic cost, etc...).


**Main idea:** replace expensive black-box queries with **cheaper computations** that will save additional queries in the long run.



**Bayesian optimization** aims to efficiently optimize **black-box** functions:


$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

**No gradients**, observations may be **corrupted by noise**.




Black-box queries are **very expensive** (time, economic cost, etc...).

**Main idea:** replace expensive black-box queries with **cheaper computations** that will save additional queries in the long run.



Done in practice by using **Bayesian machine learning**.

The Bayesian approach to machine learning: **probabilistic models** and **inference algorithms**.



The Bayesian approach to machine learning: **probabilistic models** and **inference algorithms**.



The Bayesian approach to machine learning: **probabilistic models** and **inference algorithms**.



Implemented using **probability theory** and **Bayes' rule**.

The Bayesian approach to machine learning: **probabilistic models** and **inference algorithms**.



Implemented using **probability theory** and **Bayes' rule**.

$$p(\theta | \text{Data}) = \frac{p(\text{Data} | \theta) p(\theta)}{p(\text{Data})}, \quad \text{Model}$$

The Bayesian approach to machine learning: **probabilistic models** and **inference algorithms**.



Implemented using **probability theory** and **Bayes' rule**.


$$p(\theta | \text{Data}) = \frac{p(\text{Data} | \theta) p(\theta)}{p(\text{Data})}, \quad \text{Model}$$

$$p(y | \text{Data}) = \int p(y | \theta) p(\theta | \text{Data}) d\theta. \quad \text{Inference}$$

**Gaussian processes** are the default **model** in Bayesian optimization.

### Advantages:

- They are non-parametric.
- **Inference** is exact: closed-form **Gaussian predictive distribution**.




A **covariance function**  $C(\mathbf{x}_1, \mathbf{x}_2)$  is used to measure data point similarity.



Drugs



Materials



Robots



Hardware



Optimal design in engineering


problems  $\downarrow \uparrow$  solutions

**Bayesian optimization**

Drugs



Materials



Robots




Hardware



Optimal design in engineering


problems  $\downarrow$   $\uparrow$  solutions



Drugs



Materials



Robots





Hardware




Optimal design in engineering

problems  $\downarrow$   $\uparrow$  solutions











① Get initial sample.



① Get initial sample.




① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$






① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$




① Get initial sample.

② **Fit a model to the data:**

$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$



① Get initial sample.


② Fit a model to the data:



$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .





① Get initial sample.

② Fit a model to the data:


$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:


$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ **Optimize acquisition function  $\alpha(\mathbf{x})$ .**

⑤ Collect data and update model.



- 1 Get initial sample.
- 2 Fit a model to the data:
$$p(y|\mathbf{x}, \mathcal{D}_n).$$
- 3 Select data collection strategy:
$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$
- 4 Optimize acquisition function  $\alpha(\mathbf{x})$ .
- 5 **Collect data and update model.**
- 6 Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:


$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!





① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



1 Get initial sample.

2 Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


3 Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

4 Optimize acquisition function  $\alpha(\mathbf{x})$ .

5 Collect data and update model.

6 Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$



③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:


$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!





① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$


③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .

⑤ Collect data and update model.

⑥ Repeat!



① Get initial sample.

② Fit a model to the data:

$$p(y|\mathbf{x}, \mathcal{D}_n).$$

③ Select data collection strategy:

$$\alpha(\mathbf{x}) = \mathbf{E}_{p(y|\mathbf{x}, \mathcal{D}_n)}[U(y|\mathbf{x}, \mathcal{D}_n)].$$

④ Optimize acquisition function  $\alpha(\mathbf{x})$ .


⑤ Collect data and update model.

⑥ Repeat!

# Discovering new optimal molecules

## Library generation


### Fragments




### Bonding rules




## Performance evaluation



## Interesting molecules




# Discovering new optimal molecules



Bayesian optimization can **accelerate** the search!


# Discovering new optimal molecules



Bayesian optimization can **accelerate** the search!

Need **model** to map molecules to properties:


# Discovering new optimal molecules



Bayesian optimization can **accelerate** the search!

Need **model** to map molecules to properties: **deep neural networks**.

# Discovering new optimal molecules




Bayesian optimization can **accelerate** the search!

Need **model** to map molecules to properties: **deep neural networks**.

## Challenges:


- **Probabilistic inference** is very difficult!
- Very complicated models, very large number of parameters.





$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation



$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation

$$q(\theta) =$$



**Minimize** a divergence between the posterior  $p$  and  $q$ :  $D_{\alpha}(p||q)$ .

$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation

$$q(\theta) =$$


**Minimize** a divergence between the posterior  $p$  and  $q$ :  $D_\alpha(p||q)$ .




$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation


$$q(\theta) =$$


**Minimize** a divergence between the posterior  $p$  and  $q$ :  $D_\alpha(p||q)$ .




$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation



**Minimize** a divergence between the posterior  $p$  and  $q$ :  $D_{\alpha}(p||q)$ .




$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation

$$q(\theta) = \text{[Contour Plot of } q(\theta) \text{]}$$

**Minimize** a divergence between the posterior  $p$  and  $q$ :  $D_\alpha(p||q)$ .




$$\min_q -\log Z_q - \sum_{n=1}^N \frac{1}{\alpha} \log \mathbb{E}_q \left[ \left( \frac{f_n(\theta)}{\tilde{f}(\theta)} \right)^\alpha \right]$$

$$p(y|\text{Data}) = \int p(y|\theta)p(\theta|\text{Data})d\theta \quad \text{Inference algorithm}$$

Approximation

$$q(\theta) =$$


Minimize a divergence between the posterior  $p$  and  $q$ :  $D_\alpha(p||q)$ .



Variational  
Bayes


Expectation  
propagation

$$\min_q -\log Z_q - \sum_{n=1}^N \frac{1}{\alpha} \log \mathbb{E}_q \left[ \left( \frac{f_n(\theta)}{\tilde{f}(\theta)} \right)^\alpha \right]$$


Easy to apply to complex models using **stochastic optimization** and **autodiff**.





# Example





# Example





# Example





# Example





# Example





# Example



# Example





# Example









# Example





# Example





# Example





# Example





# Example





# Example



# Example




# Example






# Example


Alpha = 0.05



Predictive distribution




Example with 2D Gaussians




# Example


Alpha = 0.13





Predictive distribution





Example with 2D Gaussians





# Example





# Example





# Example





# Example



# Example




# Example






# Example


Alpha = 0.67





Predictive distribution





Example with 2D Gaussians



# Example




# Example




# Example


Alpha = 0.90



Predictive distribution



Example with 2D Gaussians




# Example


Alpha = 0.97



Predictive distribution




Example with 2D Gaussians




# Example


Alpha = 1.00




Predictive distribution



Example with 2D Gaussians



Harvard clean energy project. Data from 60.000 **organic photovoltaics**.



Bayesian deep neural networks (BDNN), 2 hidden layers with 400 units.

**Table:** Avg. Test Error and Test Log-likelihood.


|                | $\alpha=1.0$ (EP) | $\alpha=0.5$                       | $\alpha=0$ (VB)  | Sparse GP        |
|----------------|-------------------|------------------------------------|------------------|------------------|
| Error          | $1.28 \pm 0.01$   | <b><math>1.08 \pm 0.01</math></b>  | $1.13 \pm 0.01$  | $1.37 \pm 0.01$  |
| Log-likelihood | $-0.93 \pm 0.01$  | <b><math>-0.74 \pm 0.01</math></b> | $-1.39 \pm 0.03$ | $-1.23 \pm 0.01$ |

# Take home messages

## Alpha divergence minimization...

- ① generalizes existing inference methods (VB and EP).
- ② better than any of them in regression problems with neural nets.





How to represent molecules? We use the SMILES specification.




Molecules are encoded as short **character strings**.

How to represent molecules? We use the SMILES specification.

Molecules are encoded as short **character strings**.

Graph

SMILES string

|                                                                                   |                                                                                |
|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
|  | <chem>CCC[C@@H](O)CC\C=C=C\C#CC#C\C=C\CO</chem>                                |
|  | <chem>COC(=O)C\C=C=C\C1C(C)(C)[C@H]1C(=O)O[C@@H]2C(C)=C(C(=O)C2)CC=CC=C</chem> |
|  | <chem>O1C=C[C@H]([C@H]1O2)c3c2cc(OC)c4c3OC(=O)C5=C4CCC(=O)5</chem>             |
|  | <chem>OC[C@@H](O1)[C@@H](O)[C@H](O)[C@@H](O)[C@@H](O)1</chem>                  |

Optimization over molecules is **challenging**. Existing methods either use

- Exhaustive search through a fixed library.
- Genetic algorithms.

These require manual specification.

Optimization over molecules is **challenging**. Existing methods either use

- Exhaustive search through a fixed library.
- Genetic algorithms.

These require manual specification.

**A continuous, data-driven representation** allows us to use **gradient based optimization** and can **leverage chemical databases** to avoid manual tuning.


Optimization over molecules is **challenging**. Existing methods either use

- Exhaustive search through a fixed library.
- Genetic algorithms.


These require manual specification.

**A continuous, data-driven representation** allows us to use **gradient based optimization** and can **leverage chemical databases** to avoid manual tuning.


Implemented by using seq2seq variational autoencoder with RNNs on SMILES.



# Latent space representation of 220,000 drug-like molecules




## Interpolation in latent space







Easy to add a **surrogate model** from latent space to property.



Easy to add a **surrogate model** from latent space to property.



**Gradient-based** optimization can be used in latent space.



**Data:** 250,000 molecules from Zinc database. **Surrogate model:** sparse GP.

We aim to maximize the water-octanol partition coefficient (**logP**):

$$J(m) = \log P(m) - \text{SA}(m) - \text{ring-penalty}(m),$$


We perform 10 iterations of BO, collecting 50 points per iteration.

**Data:** 250,000 molecules from Zinc database. **Surrogate model:** sparse GP.


We aim to maximize the water-octanol partition coefficient (**logP**):

$$J(m) = \log P(m) - \text{SA}(m) - \text{ring-penalty}(m),$$


We perform 10 iterations of BO, collecting 50 points per iteration.



Molecule 1



Molecule 2




Found two molecules better than anything seen before.

## Local optimization of OLED molecules



We optimize the delayed fluorescence decay rate  $\kappa_{\text{TADF}}$ , as estimated from TDDFT computations on 150,000 molecules.



# Take home messages

## Alpha divergence minimization...

- ① generalizes existing inference methods (VB and EP).
- ② better than any of them in regression problems with neural nets.

## Data-driven molecule optimization...

- ① New compounds generated by optimizing in latent space.
- ② Can leverage large lists of unlabeled chemical compounds.

# Collaborators



R. P. Adams



J. Aguilera-Ipagirre



A. Aspuru-Guzik



T. B. Bui



D. Duvenaud



R. Gomez-Bonbardelli



D. Hernández-Lobato



T. Hirzel



Y. Li



M. Rowland



R. E. Turner

Thanks!