

---

# **pyAB Documentation**

***Release 0.0.1***

**Aditya Varma Kalidindi**

**Jun 17, 2020**



---

## Contents

---

<b>1</b>	<b>Features:</b>	<b>3</b>
1.1	Installation . . . . .	3
1.2	Quick Start . . . . .	4
1.3	pyAB API . . . . .	8
1.4	Contributing to pyAB . . . . .	10
1.5	Release History . . . . .	10
<b>2</b>	<b>Usage:</b>	<b>11</b>
<b>3</b>	<b>Bayesian A/B Test</b>	<b>13</b>
<b>4</b>	<b>Frequentist A/B Test</b>	<b>15</b>
	<b>Index</b>	<b>19</b>



pyAB is a Python package for Bayesian & Frequentist A/B Testing.



# CHAPTER 1

---

## Features:

---

### Bayesian A/B Test:

- Conduct quick experiments to check for winning variant with additional prior information (Beta Distribution parameters).
- Try different evaluation metrics (Uplift Ratio, Uplift Difference & Uplift Percent Gain) & vary number of mcmc simulations.
- Visualize & inspect Uplift Density & Cumulative Density distributions.

### Frequentist A/B Test:

- Conduct quick experiments to check for winning variant using two sample proportion test (Statistical significance).
- Estimate required sample size per variant to reach provided Type-II error rate.
- Visualize & inspect power curve for varying alternative proportions.

## 1.1 Installation

Best way to install pyAB is through pip

```
pip install pyAB
```

To install from source, use the following Github link

```
git clone https://github.com/AdiVarma27/pyAB.git
cd pyAB
python setup.py install
```

### 1.1.1 Dependencies

pyAB has the following dependencies:

- numpy
- matplotlib
- seaborn
- scipy
- statsmodels

## 1.2 Quick Start

### 1.2.1 Bayesian A/B Test

Let us assume we have two Banner Ads and want to run an AB Test to decide on the final version. We run the test and collect 1000 samples per version. We observe 100 and 120 clicks for version-A & Version-B respectively (**10 % & 12.5 % Click-through-rates**). From our previous experience, we know that the average Click-through-rate for our previous Ads was around 12 %.

We first need to import ABTestBayesian class and provide prior clicks success\_prior and prior impressions trials\_prior. Then, call the conduct\_experiment method with successful clicks and impressions per version.

For uplift\_method, there are three metrics to choose from are 'uplift\_ratio', 'uplift\_percent' & 'uplift\_difference'. We also choose mcmc num\_simulations, which samples from Uplift Probability Density function.

```
# import Bayesian class
from pyab.experiments import ABTestBayesian

# provide beta priors
ad_experiment_bayesian = ABTestBayesian(success_prior=120, trials_prior=1000)

# conduct experiment with two variants successes and trials, along with uplift method
↳and number of simulations
ad_experiment_bayesian.conduct_experiment(success_null=100, trials_null=1000,
                                          success_alt=125, trials_alt=1000,
                                          uplift_method='uplift_ratio', num_
↳simulations=1000)
```

Bayesian A/B test results can extremely useful to **understand & communicate test results** with other stakeholders and answers the main business question: **Which version works the best ?**

#### Output:

```
pyAB Summary
=====

Test Parameters
-----

Variant A: Successful Trials 100, Sample Size 1000
Variant B: Successful Trials 125, Sample Size 1000
Prior: Successful Trials 120, Sample Size 1000

Test Results
-----
```

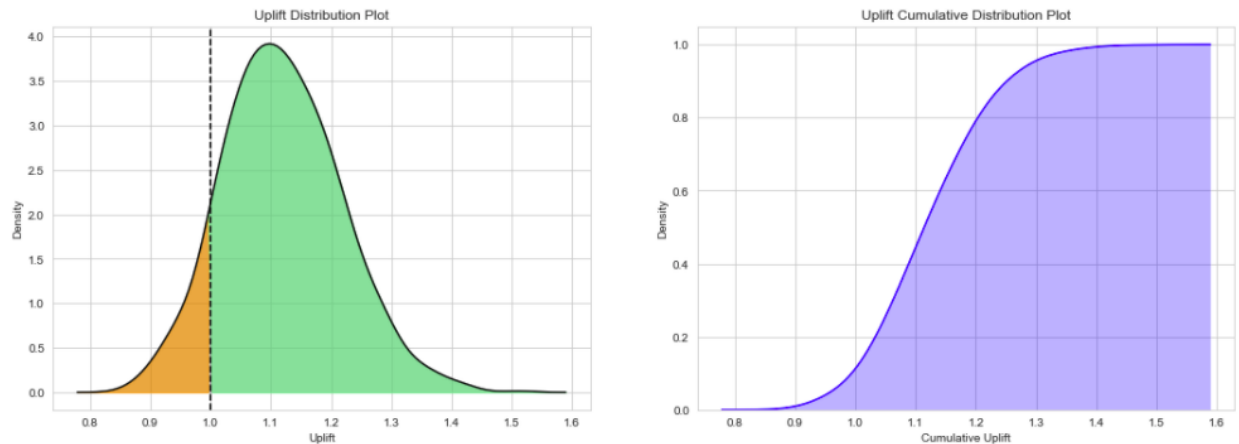
(continues on next page)



(continued from previous page)

```
Evaluation Metric: uplift_ratio
Number of mcmc simulations: 1000

90.33 % simulations show Uplift Ratio above 1.
```



## 1.2.2 Frequentist A/B Test

Let us now run a Frequentist A/B Test and verify if there is a statistically significant difference between two proportions, provided the sample sizes and Type-I Error rate. From above, we know the performance of version-A & version-B (**10 % & 12.5 % Click-through-rates**), for 1000 impressions of each version.

We first need to import `ABTestFrequentist` class and provide type of alternative hypothesis `alt_hypothesis`, 'one\_tailed' or 'two\_tailed' & Type-I error rate `alpha` (default = 0.05). Then, we call the `conduct_experiment` method with successful clicks and impressions per version.

This traditional methodology might be **slightly tricky to communicate**, and **Type-I & Type-II error rates** need to be accounted for, unlike Bayesian methods.

```
# import Frequentist class
from pyab.experiments import ABTestFrequentist

# provide significance rate and type of test
ad_experiment_freq = ABTestFrequentist(alpha=0.05, alt_hypothesis='one_tailed')

# conduct experiment with two variants successes and trials, returns stat & pvalue
stat, pvalue = ad_experiment_freq.conduct_experiment(success_null=100, trials_
↳ null=1000,
                                     success_alt=125, trials_alt=1000)
```

### Output:

```
pyAB Summary
=====

Test Parameters
_____
```

(continues on next page)


(continued from previous page)

```
Variant A: Success Rate 0.1, Sample Size 1000
Variant B: Success Rate 0.125, Sample Size 1000
Type-I Error: 0.05, one_tailed test
```

#### Test Results

```
Test Stat: 1.769
p-value: 0.038
Type-II Error: 0.451
Power: 0.549
```

There **is** a statistically significant difference **in** proportions of two variants.







Given that the current Type-II error is 0.451 at 1000 samples per variant, we can find out **required sample size per variant** to reach Type-II error of 0.1.

```
# required sample size per variant for given beta
ad_experiment.get_sample_size(beta=0.1)
```

#### Output:

```
2729
```

**Never misinterpret your Results !**



## 1.3 pyAB API

(Version 0.0.1)

### 1.3.1 Experiments

**class** `pyab.experiments.ABTestBayesian` (*success\_prior*, *trials\_prior*)  
Bayesian A/B Testing.

#### Parameters

- **success\_prior** (*int*) – Number of successful samples from prior.
- **trials\_prior** (*int*) – Number of trials from prior.

**calculate\_uplift\_area** ()  
Calculate Uplift pdf & area beyond threshold.

#### Returns

- **uplift\_distribution** (*ndarray*) – uplift distribution based on chosen uplift method.
- **uplift\_area** (*float*) – percentage area above threshold.

**conduct\_experiment** (*success\_null*, *trials\_null*, *success\_alt*, *trials\_alt*, *uplift\_method*=*'uplift\_percent'*, *num\_simulations*=1000)  
Conduct experiment & generate uplift distributions.

#### Parameters

- **success\_null** (*int*) – Number of successful samples for variant-a.
- **trials\_null** (*int*) – Number of trials for variant-a.
- **success\_alt** (*int*) – Number of successful samples for variant-b.
- **trials\_alt** (*int*) – Number of trials for variant-b.
- **num\_simulations** (*int*) – Number of mcmc simulations.
- **uplift\_method** (*str*, *default* = *'uplift\_percent'*) – Uplift evaluation metric.
  - **'uplift\_percent'**: percent uplift gain from variant-a to variant-b
  - **'uplift\_ratio'**: uplift ratio of variant-b & variant-a
  - **'uplift\_difference'**: uplift difference between variant-b & variant-a

**plot\_uplift\_distributions** (*figsize*=(18, 6))  
Plot uplift pdf & cdf for provided experiment parameters.

Parameters **figsize** (*tuple*, *default* = (18, 6)) – matplotlib plot size.

**print\_bayesian\_results** ()  
Print Bayesian Experiment Results

**class** `pyab.experiments.ABTestFrequentist` (*alpha*=0.05, *alt\_hypothesis*=*'one\_tailed'*)  
Frequentist A/B Testing aka Two sample proportion test.

#### Parameters

- **alpha** (*float*, *default* = 0.05) – Significance level or Type 1 error rate.

- **alt\_hypothesis** (*str*, *default* = 'one\_tailed') – One or two tailed hypothesis test.
  - 'one\_tailed': one tailed hypothesis test
  - 'two\_tailed': two tailed hypothesis test

**calculate\_power** (*stat*)

Calculate power (1-beta) at given test statistics.

**Parameters** **stat** (*float*) – z or t test statistic.

**Returns** **1 - beta** – power at given test statistic.

**Return type** float

**calculate\_stat** (*prop\_alt*)

Calculate test statistic with current experiment parameters.

**Parameters** **prop\_alt** (*float*) – alternate hypothesis proportion.

**Returns** **stat** – z or t statistic.

**Return type** float

**conduct\_experiment** (*success\_null*, *trials\_null*, *success\_alt*, *trials\_alt*)

Conduct experiment & generate power curve with provided parameters.

**Parameters**

- **success\_null** (*int*) – number of successful clicks or successful events (Version-A).
- **trials\_null** (*int*) – number of impressions or events (Version-A).
- **success\_alt** (*int*) – number of successful clicks or successful events (Version-B).
- **trials\_alt** (*int*) – number of impressions or events (Version-B).

**Returns**

- **stat** (*float*) – z or t statistic.
- **pvalue** (*float*) – probability of obtaining results atleast as extreme as the results actually observed during the test.

**get\_sample\_size** (*beta=0.1*)

Calculate required sample size per group to obtain provided beta.

**Parameters** **beta** (*float*) – Type 2 error rate.

**Returns** **n** – sample size per group.

**Return type** int

**plot\_power\_curve** (*figsize=(9, 6)*)

Plot power curve for provided experiment parameters.

**Parameters** **figsize** (*tuple*, *default* = (9, 6)) – matplotlib plot size.

**print\_freq\_results** ()

Print Frequentist Experiment Results

## 1.4 Contributing to pyAB

Welcome! pyAB is a community project and your contribution is important to the packages usability and success.

### 1.4.1 Code of Conduct

Contributors and participants of pyAB are expected to follow guidelines provided by [Python Community Code of Conduct](#).

### 1.4.2 General Guidelines

- For issues, please submit them to the [issue tracker](#). If you can provide features, improvements or anything else the world of AB Testing has to offer, your contributions are highly appreciated.
- Code in master branch should reflect the latest version. Create a pull request and add your merge request to the `dev-pyab` branch.
- Please follow [pep8](#) for coding conventions. For quick information about Git, visit <https://rogerdudler.github.io/git-guide/>.

## 1.5 Release History

### 1.5.1 Version 0.0.1

- ABTestBayesian, ABTestFrequentist classes in experiments module.
- Base utility functions, test cases.

## CHAPTER 2

---

Usage:

---





## CHAPTER 3

---

### Bayesian A/B Test

---

Let us assume we have two Banner Ads and want to run an AB Test to decide on the final version. We run the test and collect 1000 samples per version. We observe 100 and 120 clicks for version-A & Version-B respectively (**10 % & 12.5 % Click-through-rates**). From our previous experience, we know that the average Click-through-rate for our previous Ads was around 12 %.

We first need to import ABTestBayesian class and provide prior clicks success\_prior and prior impressions trials\_prior. Then, call the conduct\_experiment method with successful clicks and impressions per version.

For uplift\_method, there are three metrics to choose from are 'uplift\_ratio', 'uplift\_percent' & 'uplift\_difference'. We also choose mcmc num\_simulations, which samples from Uplift Probability Density function.

```
# import Bayesian class
from pyab.experiments import ABTestBayesian

# provide beta priors
ad_experiment_bayesian = ABTestBayesian(success_prior=120, trials_prior=1000)

# conduct experiment with two variants successes and trials, along with uplift method_
↪and number of simulations
ad_experiment_bayesian.conduct_experiment(success_null=100, trials_null=1000,
                                           success_alt=125, trials_alt=1000,
                                           uplift_method='uplift_ratio', num_
↪simulations=1000)
```

Bayesian A/B test results can extremely useful to **understand & communicate test results** with other stakeholders and answers the main business question: **Which version works the best ?**

#### Output:

```
pyAB Summary
=====

Test Parameters
```

(continues on next page)

(continued from previous page)

---

```
Variant A: Successful Trials 100, Sample Size 1000
Variant B: Successful Trials 125, Sample Size 1000
Prior: Successful Trials 120, Sample Size 1000
```



---

#### Test Results

---

```
Evaluation Metric: uplift_ratio
Number of mcmc simulations: 1000
```

```
90.33 % simulations show Uplift Ratio above 1.
```



## CHAPTER 4

---

### Frequentist A/B Test

---

Let us now run a Frequentist A/B Test and verify if there is a significant difference between two proportions provided the sample sizes and Type-I Error rate. From above, we know the performance of version-A & version-B (**10 % & 12.5 % Click-through-rates**), for 1000 impressions of each version.

We first need to import ABTestFrequentist class and provide type of alternative hypothesis `alt_hypothesis`, 'one\_tailed' or 'two\_tailed' & Type-I error rate `alpha` (default = 0.05). Then, we call the `conduct_experiment` method with successful clicks and impressions per version.

This traditional methodology might be **slightly tricky to communicate**, and **Type-I & Type-II error rates** need to be accounted for, unlike Bayesian methods.

```
# import Frequentist class
from pyab.experiments import ABTestFrequentist

# provide significance rate and type of test
ad_experiment_freq = ABTestFrequentist(alpha=0.05, alt_hypothesis='one_tailed')

# conduct experiment with two variants successes and trials, returns stat & pvalue
stat, pvalue = ad_experiment_freq.conduct_experiment(success_null=100, trials_
↳ null=1000,
                                     success_alt=125, trials_alt=1000)
```

#### Output:

```
pyAB Summary
=====

Test Parameters
-----

Variant A: Success Rate 0.1, Sample Size 1000
Variant B: Success Rate 0.125, Sample Size 1000
Type-I Error: 0.05, one_tailed test
```


(continues on next page)

(continued from previous page)

### Test Results

Test Stat: 1.769  
p-value: 0.038  
Type-II Error: 0.451  
Power: 0.549

There **is** a statistically significant difference **in** proportions of two variants.




Given that the current Type-II error is 0.451 at 1000 samples per variant, we can find out **required sample size per variant** to reach Type-II error of 0.1.

```
# required sample size per variant for given beta  
ad_experiment.get_sample_size(beta=0.1)
```

### Output:

```
2729
```

**Never misinterpret your Results !**





## A

ABTestBayesian (*class in pyab.experiments*), 8  
ABTestFrequentist (*class in pyab.experiments*), 8

## C

calculate\_power()  
    (*pyab.experiments.ABTestFrequentist method*),  
    9  
calculate\_stat() (*pyab.experiments.ABTestFrequentist  
    method*), 9  
calculate\_uplift\_area()  
    (*pyab.experiments.ABTestBayesian method*), 8  
conduct\_experiment()  
    (*pyab.experiments.ABTestBayesian method*), 8  
conduct\_experiment()  
    (*pyab.experiments.ABTestFrequentist method*),  
    9

## G

get\_sample\_size()  
    (*pyab.experiments.ABTestFrequentist method*),  
    9

## P

plot\_power\_curve()  
    (*pyab.experiments.ABTestFrequentist method*),  
    9  
plot\_uplift\_distributions()  
    (*pyab.experiments.ABTestBayesian method*), 8  
print\_bayesian\_results()  
    (*pyab.experiments.ABTestBayesian method*), 8  
print\_freq\_results()  
    (*pyab.experiments.ABTestFrequentist method*),  
    9