

---

# Numpyro Documentation

*Release 0.0*

**Uber AI Labs**

**Jun 03, 2019**



<b>1</b>	<b>Markov Chain Monte Carlo (MCMC)</b>	<b>1</b>
1.1	Hamiltonian Monte Carlo . . . . .	1
1.2	MCMC Utilities . . . . .	4
<b>2</b>	<b>Stochastic Variational Inference (SVI)</b>	<b>7</b>
2.1	ELBo . . . . .	8
<b>3</b>	<b>Base Distribution</b>	<b>9</b>
3.1	Distribution . . . . .	9
3.2	TransformedDistribution . . . . .	10
<b>4</b>	<b>Continuous Distributions</b>	<b>13</b>
4.1	Beta . . . . .	13
4.2	Cauchy . . . . .	13
4.3	Chi2 . . . . .	14
4.4	Dirichlet . . . . .	14
4.5	Exponential . . . . .	14
4.6	Gamma . . . . .	15
4.7	GaussianRandomWalk . . . . .	15
4.8	HalfCauchy . . . . .	15
4.9	HalfNormal . . . . .	16
4.10	LKJCholesky . . . . .	16
4.11	LogNormal . . . . .	17
4.12	Normal . . . . .	17
4.13	Pareto . . . . .	17
4.14	StudentT . . . . .	18
4.15	TruncatedCauchy . . . . .	18
4.16	TruncatedNormal . . . . .	18
4.17	Uniform . . . . .	19
<b>5</b>	<b>Discrete Distributions</b>	<b>21</b>
5.1	Bernoulli . . . . .	21
5.2	BernoulliLogits . . . . .	21
5.3	BernoulliProbs . . . . .	21
5.4	Binomial . . . . .	22
5.5	BinomialLogits . . . . .	22
5.6	BinomialProbs . . . . .	22

---

5.7	Categorical . . . . .	23
5.8	CategoricalLogits . . . . .	23
5.9	CategoricalProbs . . . . .	23
5.10	Multinomial . . . . .	24
5.11	MultinomialLogits . . . . .	24
5.12	MultinomialProbs . . . . .	24
5.13	Poisson . . . . .	24
<b>6</b>	<b>Pyro Primitives</b>	<b>27</b>
<b>7</b>	<b>Effect Handlers</b>	<b>29</b>
<b>8</b>	<b>Autocorrelation</b>	<b>33</b>
<b>9</b>	<b>Autocovariance</b>	<b>35</b>
<b>10</b>	<b>Effective Sample Size</b>	<b>37</b>
<b>11</b>	<b>Gelman Rubin</b>	<b>39</b>
<b>12</b>	<b>Split Gelman Rubin</b>	<b>41</b>
<b>13</b>	<b>HPDI</b>	<b>43</b>
<b>14</b>	<b>Summary</b>	<b>45</b>
<b>15</b>	<b>Indices and tables</b>	<b>47</b>
	<b>Python Module Index</b>	<b>49</b>
	<b>Index</b>	<b>51</b>

# CHAPTER 1

---

## Markov Chain Monte Carlo (MCMC)

---

### 1.1 Hamiltonian Monte Carlo

`mcmc` (*num\_warmup*, *num\_samples*, *init\_params*, *sampler='hmc'*, *constrain\_fn=None*, *print\_summary=True*,  
    `**sampler_kwargs`)

Convenience wrapper for MCMC samplers – runs warmup, prints diagnostic summary and returns a collections of samples from the posterior.

#### Parameters

- **num\_warmup** – Number of warmup steps.
- **num\_samples** – Number of samples to generate from the Markov chain.
- **init\_params** – Initial parameters to begin sampling. The type can must be consistent with the input type to *potential\_fn*.
- **sampler** – currently, only *hmc* is implemented (default).
- **constrain\_fn** – Callable that converts a collection of unconstrained sample values returned from the sampler to constrained values that lie within the support of the sample sites.
- **print\_summary** – Whether to print diagnostics summary for each sample site. Default is True.
- **\*\*sampler\_kwargs** – Sampler specific keyword arguments.
  - *HMC*: Refer to `hmc()` and `init_kernel()` for accepted arguments. Note that all arguments must be provided as keywords.

**Returns** collection of samples from the posterior.

```
>>> true_coefs = np.array([1., 2., 3.])
>>> data = random.normal(random.PRNGKey(2), (2000, 3))
>>> dim = 3
>>> labels = dist.Bernoulli(logits=(true_coefs * data).sum(-1)).sample(random.
    ↵PRNGKey(3))
```

(continues on next page)

(continued from previous page)

```

>>>
>>> def model(data, labels):
...     coefs_mean = np.zeros(dim)
...     coefs = sample('beta', dist.Normal(coefs_mean, np.ones(3)))
...     intercept = sample('intercept', dist.Normal(0., 10.))
...     return sample('y', dist.Bernoulli(logits=(coefs * data + intercept).sum(-
... 1)), obs=labels)
>>>
>>> init_params, potential_fn, constrain_fn = initialize_model(random.PRNGKey(0),
... model,
...                                         data, labels)
...
>>> num_warmup, num_samples = 1000, 1000
>>> samples = mcmc(num_warmup, num_samples, init_params,
...                  potential_fn=potential_fn,
...                  constrain_fn=constrain_fn)
warmup: 100%|| 1000/1000 [00:09<00:00, 109.40it/s, 1 steps of size 5.83e-01. acc.=
prob=0.79]
sample: 100%|| 1000/1000 [00:00<00:00, 1252.39it/s, 1 steps of size 5.83e-01. acc.=
prob=0.85]


```

	mean	sd	5.5%	94.5%	n_eff	Rhat
coefs[0]	0.96	0.07	0.85	1.07	455.35	1.01
coefs[1]	2.05	0.09	1.91	2.20	332.00	1.01
coefs[2]	3.18	0.13	2.96	3.37	320.27	1.00
intercept	-0.03	0.02	-0.06	0.00	402.53	1.00

**hmc**(*potential\_fn*, *kinetic\_fn=None*, *algo='NUTS'*)

Hamiltonian Monte Carlo inference, using either fixed number of steps or the No U-Turn Sampler (NUTS) with adaptive path length.

**References:**

1. *MCMC Using Hamiltonian Dynamics*, Radford M. Neal
2. *The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo*, Matthew D. Hoffman, and Andrew Gelman.
3. *A Conceptual Introduction to Hamiltonian Monte Carlo*, Michael Betancourt

**Parameters**

- ***potential\_fn*** – Python callable that computes the potential energy given input parameters. The input parameters to *potential\_fn* can be any python collection type, provided that *init\_params* argument to *init\_kernel* has the same type.
- ***kinetic\_fn*** – Python callable that returns the kinetic energy given inverse mass matrix and momentum. If not provided, the default is euclidean kinetic energy.
- ***algo* (*str*)** – Whether to run HMC with fixed number of steps or NUTS with adaptive path length. Default is NUTS.

**Returns** a tuple of callables (*init\_kernel*, *sample\_kernel*), the first one to initialize the sampler, and the second one to generate samples given an existing one.

**Example**

```

>>> true_coefs = np.array([1., 2., 3.])
>>> data = random.normal(random.PRNGKey(2), (2000, 3))
>>> dim = 3
>>> labels = dist.Bernoulli(logits=(true_coefs * data).sum(-1)).sample(random.
->PRNGKey(3))
>>>
>>> def model(data, labels):
...     coefs_mean = np.zeros(dim)
...     coefs = sample('beta', dist.Normal(coefs_mean, np.ones(3)))
...     intercept = sample('intercept', dist.Normal(0., 10.))
...     return sample('y', dist.Bernoulli(logits=(coefs * data + intercept).sum(-
->1)), obs=labels)
>>>
>>> init_params, potential_fn, constrain_fn = initialize_model(random.PRNGKey(0),
...
->labels)
>>> init_kernel, sample_kernel = hmc(potential_fn, algo='NUTS')
>>> hmc_state = init_kernel(init_params,
...                           trajectory_length=10,
...                           num_warmup=300)
>>> samples = fori_collect(500, sample_kernel, hmc_state,
...                         transform=lambda state: constrain_fn(state.z))
>>> print(np.mean(samples['beta'], axis=0))
[0.9153987 2.0754058 2.9621222]

```

**init\_kernel**(*init\_params*, *num\_warmup*, *step\_size*=1.0, *adapt\_step\_size*=True, *adapt\_mass\_matrix*=True, *dense\_mass*=False, *target\_accept\_prob*=0.8, *trajectory\_length*=6.283185307179586, *max\_tree\_depth*=10, *run\_warmup*=True, *progbar*=True, *rng*=DeviceArray([0, 0], *dtype*=uint32))

Initializes the HMC sampler.

#### Parameters

- **init\_params** – Initial parameters to begin sampling. The type must be consistent with the input type to *potential\_fn*.
- **num\_warmup\_steps** (*int*) – Number of warmup steps; samples generated during warmup are discarded.
- **step\_size** (*float*) – Determines the size of a single step taken by the verlet integrator while computing the trajectory using Hamiltonian dynamics. If not specified, it will be set to 1.
- **adapt\_step\_size** (*bool*) – A flag to decide if we want to adapt step\_size during warm-up phase using Dual Averaging scheme.
- **adapt\_mass\_matrix** (*bool*) – A flag to decide if we want to adapt mass matrix during warm-up phase using Welford scheme.
- **dense\_mass** (*bool*) – A flag to decide if mass matrix is dense or diagonal (default when *dense\_mass*=False)
- **target\_accept\_prob** (*float*) – Target acceptance probability for step size adaptation using Dual Averaging. Increasing this value will lead to a smaller step size, hence the sampling will be slower but more robust. Default to 0.8.
- **trajectory\_length** (*float*) – Length of a MCMC trajectory for HMC. Default value is  $2\pi$ .

- **max\_tree\_depth** (`int`) – Max depth of the binary tree created during the doubling scheme of NUTS sampler. Defaults to 10.
- **run\_warmup** (`bool`) – Flag to decide whether warmup is run. If `True`, `init_kernel` returns an initial `HMCState` that can be used to generate samples using MCMC. Else, returns the arguments and callable that does the initial adaptation.
- **progressbar** (`bool`) – Whether to enable progress bar updates. Defaults to `True`.
- **heuristic\_step\_size** (`bool`) – If `True`, a coarse grained adjustment of step size is done at the beginning of each adaptation window to achieve `target_acceptance_prob`.
- **rng** (`jax.random.PRNGKey`) – random key to be used as the source of randomness.

#### `sample_kernel(hmc_state)`

Given an existing `HMCState`, run HMC with fixed (possibly adapted) step size and return a new `HMCState`.

**Parameters** `hmc_state` – Current sample (and associated state).

**Returns** new proposed `HMCState` from simulating Hamiltonian dynamics given existing state.

`HMCState = <class 'numpyro.mcmc.HMCState'>`

A `namedtuple()` consisting of the following fields:

- **i** - iteration. This is reset to 0 after warmup.
- **z** - Python collection representing values (unconstrained samples from the posterior) at latent sites.
- **z\_grad** - Gradient of potential energy w.r.t. latent sample sites.
- **potential\_energy** - Potential energy computed at the given value of `z`.
- **num\_steps** - Number of steps in the Hamiltonian trajectory (for diagnostics).
- **accept\_prob** - Acceptance probability of the proposal. Note that `z` does not correspond to the proposal if it is rejected.
- **mean\_accept\_prob** - Mean acceptance probability until current iteration during warmup adaptation or sampling (for diagnostics).
- **step\_size** - Step size to be used by the integrator in the next iteration. This is adapted during warmup.
- **inverse\_mass\_matrix** - The inverse mass matrix to be used for the next iteration. This is adapted during warmup.
- **rng** - random number generator seed used for the iteration.

## 1.2 MCMC Utilities

### `initialize_model(rng, model, *model_args, init_strategy='uniform', **model_kwargs)`

Given a model with Pyro primitives, returns a function which, given unconstrained parameters, evaluates the potential energy (negative joint density). In addition, this also returns initial parameters sampled from the prior to initiate MCMC sampling and functions to transform unconstrained values at sample sites to constrained values within their respective support.

**Parameters**

- **rng** (`jax.random.PRNGKey`) – random number generator seed to sample from the prior.
- **model** – Python callable containing Pyro primitives.
- **\*model\_args** – args provided to the model.

- **init\_strategy** (*str*) – initialization strategy - *uniform* initializes the unconstrained parameters by drawing from a *Uniform(-2, 2)* distribution (as used by Stan), whereas *prior* initializes the parameters by sampling from the prior for each of the sample sites.
- **\*\*model\_kwargs** – kwargs provided to the model.

**Returns** tuple of (*init\_params*, *potential\_fn*, *constrain\_fn*), *init\_params* are values from the prior used to initiate MCMC, *constrain\_fn* is a callable that uses inverse transforms to convert unconstrained HMC samples to constrained values that lie within the site's support.

**fori\_collect** (*n*, *body\_fun*, *init\_val*, *transform*=<function identity>, *progbar=True*, *\*\*progbar\_opts*)

This looping construct works like `fori_loop()` but with the additional effect of collecting values from the loop body. In addition, this allows for post-processing of these samples via *transform*, and progress bar updates. Note that, in some cases, *progbar=False* can be faster, when collecting a lot of samples. Refer to example usage in `hmc()`.

#### Parameters

- **n** (*int*) – number of times to run the loop body.
- **body\_fun** – a callable that takes a collection of *np.ndarray* and returns a collection with the same shape and *dtype*.
- **init\_val** – initial value to pass as argument to *body\_fun*. Can be any Python collection type containing *np.ndarray* objects.
- **transform** – A callable
- **progbar** – whether to post progress bar updates.
- **\*\*progbar\_opts** – optional additional progress bar arguments. A *diagnostics\_fn* can be supplied which when passed the current value from *body\_fun* returns a string that is used to update the progress bar postfix. Also a *progbar\_desc* keyword argument can be supplied which is used to label the progress bar.

**Returns** collection with the same type as *init\_val* with values collected along the leading axis of *np.ndarray* objects.

**summary** (*samples*, *prob=0.89*)

Prints a summary table displaying diagnostics of *samples* from the posterior. The diagnostics displayed are mean, standard deviation, the 89% Credibility Interval, `effective_sample_size()`, `split_gelman_rubin()`.

#### Parameters

- **samples** – a collection of input samples.
- **prob** (*float*) – the probability mass of samples within the HPDI interval.



# CHAPTER 2

---

## Stochastic Variational Inference (SVI)

---

**svi** (*model, guide, loss, optim\_init, optim\_update, get\_params, \*\*kwargs*)

Stochastic Variational Inference given an ELBo loss objective.

### Parameters

- **model** – Python callable with Pyro primitives for the model.
- **guide** – Python callable with Pyro primitives for the guide (recognition network).
- **loss** – ELBo loss, i.e. negative Evidence Lower Bound, to minimize.
- **optim\_init** – initialization function returned by a JAX optimizer. see: `jax.experimental.optimizers`.
- **optim\_update** – update function for the optimizer
- **get\_params** – function to get current parameters values given the optimizer state.
- **\*\*kwargs** – static arguments for the model / guide, i.e. arguments that remain constant during fitting.

**Returns** tuple of (*init\_fn, update\_fn, evaluate*).

**init\_fn** (*rng, model\_args=(), guide\_args=(), params=None*)

### Parameters

- **rng** (`jax.random.PRNGKey`) – random number generator seed.
- **model\_args** (`tuple`) – arguments to the model (these can possibly vary during the course of fitting).
- **guide\_args** (`tuple`) – arguments to the guide (these can possibly vary during the course of fitting).
- **params** (`dict`) – initial parameter values to condition on. This can be useful for

**Returns** initial optimizer state.

**update\_fn** (*i, opt\_state, rng, model\_args=(), guide\_args=()*)

Take a single step of SVI (possibly on a batch / minibatch of data), using the optimizer.

### Parameters

- **i** (`int`) – represents the i'th iteration over the epoch, passed as an argument to the optimizer's update function.
- **opt\_state** – current optimizer state.
- **rng** (`jax.random.PRNGKey`) – random number generator seed.
- **model\_args** (`tuple`) – dynamic arguments to the model.
- **guide\_args** (`tuple`) – dynamic arguments to the guide.

**Returns** tuple of (`loss_val`, `opt_state`, `rng`).

**evaluate** (`opt_state`, `rng`, `model_args=()`, `guide_args=()`)

Take a single step of SVI (possibly on a batch / minibatch of data).

### Parameters

- **opt\_state** – current optimizer state.
- **rng** (`jax.random.PRNGKey`) – random number generator seed.
- **model\_args** (`tuple`) – arguments to the model (these can possibly vary during the course of fitting).
- **guide\_args** (`tuple`) – arguments to the guide (these can possibly vary during the course of fitting).

**Returns** evaluate ELBo loss given the current parameter values (held within `opt_state`).

## 2.1 ELBo

**elbo** (`param_map`, `model`, `guide`, `model_args`, `guide_args`, `kwargs`)

This is the most basic implementation of the Evidence Lower Bound, which is the fundamental objective in Variational Inference. This implementation has various limitations (for example it only supports random variables with reparameterized samplers) but can be used as a template to build more sophisticated loss objectives.

For more details, refer to [http://pyro.ai/examples/svi\\_part\\_i.html](http://pyro.ai/examples/svi_part_i.html).

### Parameters

- **param\_map** (`dict`) – dictionary of current parameter values keyed by site name.
- **model** – Python callable with Pyro primitives for the model.
- **guide** – Python callable with Pyro primitives for the guide (recognition network).
- **model\_args** (`tuple`) – arguments to the model (these can possibly vary during the course of fitting).
- **guide\_args** (`tuple`) – arguments to the guide (these can possibly vary during the course of fitting).
- **kwargs** (`dict`) – static keyword arguments to the model / guide.

**Returns** negative of the Evidence Lower Bound (ELBo) to be minimized.

# CHAPTER 3

---

## Base Distribution

---

### 3.1 Distribution

```
class Distribution(batch_shape=(), event_shape=(), validate_args=None)
```

Bases: `object`

Base class for probability distributions in NumPyro. The design largely follows from `torch.distributions`.

#### Parameters

- **batch\_shape** – The batch shape for the distribution. This designates independent (possibly non-identical) dimensions of a sample from the distribution. This is fixed for a distribution instance and is inferred from the shape of the distribution parameters.
- **event\_shape** – The event shape for the distribution. This designates the dependent dimensions of a sample from the distribution. These are collapsed when we evaluate the log probability density of a batch of samples using `.log_prob`.
- **validate\_args** – Whether to enable validation of distribution parameters and arguments to `.log_prob` method.

As an example:

```
>>> d = dist.Dirichlet(np.ones((2, 3, 4)))
>>> d.batch_shape
(2, 3)
>>> d.event_shape
(4,)
```

```
arg_constraints = {}
support = None
reparametrized_params = []
```

**batch\_shape**

Returns the shape over which the distribution parameters are batched.

**Returns** batch shape of the distribution.

**Return type** tuple

**event\_shape**

Returns the shape of a single sample from the distribution without batching.

**Returns** event shape of the distribution.

**Return type** tuple

**sample**(key, sample\_shape=())

Returns a sample from the distribution having shape given by *sample\_shape* + *batch\_shape* + *event\_shape*.

Note that when *sample\_shape* is non-empty, leading dimensions (of size *sample\_shape*) of the returned sample will be filled with iid draws from the distribution instance.

**Parameters**

- **key** (*jax.random.PRNGKey*) – the rng key to be used for the distribution.
- **size** – the sample shape for the distribution.

**Returns** a *numpy.ndarray* of shape *sample\_shape* + *batch\_shape* + *event\_shape*

**log\_prob**(value)

Evaluates the log probability density for a batch of samples given by *value*.

**Parameters** **value** – A batch of samples from the distribution.

**Returns** a *numpy.ndarray* with shape *value.shape[:-self.event\_shape]*

**mean**

Mean of the distribution.

**variance**

Variance of the distribution.

## 3.2 TransformedDistribution

**class TransformedDistribution**(base\_distribution, transforms, validate\_args=None)

Bases: *numpyro.distributions.distribution.Distribution*

Returns a distribution instance obtained as a result of applying a sequence of transforms to a base distribution. For an example, see LogNormal and HalfNormal.

**Parameters**

- **base\_distribution** – the base distribution over which to apply transforms.
- **transforms** – a single transform or a list of transforms.
- **validate\_args** – Whether to enable validation of distribution parameters and arguments to *.log\_prob* method.

**arg\_constraints** = {}

**support**

**is\_reparametrized**

**sample** (*key, sample\_shape=()*)

See `numpyro.distributions.distribution.Distribution.sample()`

**log\_prob** (*value*)

See `numpyro.distributions.distribution.Distribution.log_prob()`

**mean**

**variance**



# CHAPTER 4

---

## Continuous Distributions

---

### 4.1 Beta

```
class Beta(concentration1, concentration0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'concentration0': <numpyro.distributions.constraints._GreaterThan object>
support = <numpyro.distributions.constraints._Interval object>
sample(key, sample_shape=())
See numpyro.distributions.distribution.Distribution.sample()
log_prob(value)
See numpyro.distributions.distribution.Distribution.log_prob()
mean
See numpyro.distributions.distribution.Distribution.mean()
variance
See numpyro.distributions.distribution.Distribution.variance()
```

### 4.2 Cauchy

```
class Cauchy(loc=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'loc': <numpyro.distributions.constraints._Real object>, 'scale':
support = <numpyro.distributions.constraints._Real object>
reparametrized_params = ['loc', 'scale']
sample(key, sample_shape=())
See numpyro.distributions.distribution.Distribution.sample()
```

```
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log\_prob\(\)

mean
    See numpyro.distributions.distribution.Distribution.mean\(\)

variance
    See numpyro.distributions.distribution.Distribution.variance\(\)
```

## 4.3 Chi2

```
class Chi2(df, validate_args=None)
    Bases: numpyro.distributions.continuous.Gamma
    arg_constraints = {'df': <numpyro.distributions.constraints.\_GreaterThan object>}
```

## 4.4 Dirichlet

```
class Dirichlet(concentration, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'concentration': <numpyro.distributions.constraints.\_GreaterThan object>}
    support = <numpyro.distributions.constraints.\_Simplex object>
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample\(\)

    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log\_prob\(\)

    mean
        See numpyro.distributions.distribution.Distribution.mean\(\)

    variance
        See numpyro.distributions.distribution.Distribution.variance\(\)
```

## 4.5 Exponential

```
class Exponential(rate=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    reparametrized_params = ['rate']
    arg_constraints = {'rate': <numpyro.distributions.constraints.\_GreaterThan object>}
    support = <numpyro.distributions.constraints.\_GreaterThan object>
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample\(\)

    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log\_prob\(\)

    mean
        See numpyro.distributions.distribution.Distribution.mean\(\)
```

**variance**

See `numpyro.distributions.distribution.Distribution.variance()`

## 4.6 Gamma

```
class Gamma(concentration, rate=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'concentration': <numpyro.distributions.constraints._GreaterThan object>}
    support = <numpyro.distributions.constraints._GreaterThan object>
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

## 4.7 GaussianRandomWalk

```
class GaussianRandomWalk(scale=1.0, num_steps=1, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'num_steps': <numpyro.distributions.constraints._IntegerGreaterThanOrEqualTo object>}
    support = <numpyro.distributions.constraints._Real object>
    reparametrized_params = ['scale']
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

## 4.8 HalfCauchy

```
class HalfCauchy(scale=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.TransformedDistribution
    reparametrized_params = ['scale']
    arg_constraints = {'scale': <numpyro.distributions.constraints._GreaterThan object>}
```

```
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log_prob()

mean
    See numpyro.distributions.distribution.Distribution.mean()

variance
    See numpyro.distributions.distribution.Distribution.variance()
```

## 4.9 HalfNormal

```
class HalfNormal(scale=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.TransformedDistribution

    reparametrized_params = ['scale']
    arg_constraints = {'scale': <numpyro.distributions.constraints._GreaterThan object>}

    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log_prob()

    mean
        See numpyro.distributions.distribution.Distribution.mean()

    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

## 4.10 LKJCholesky

```
class LKJCholesky(dimension, concentration=1.0, sample_method='onion', validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
```

LKJ distribution for lower Cholesky factors of correlation matrices. The distribution is controlled by concentration parameter  $\eta$  to make the probability of the correlation matrix  $M$  generated from a Cholesky factor proportional to  $\det(M)^{\eta-1}$ . Because of that, when concentration == 1, we have a uniform distribution over Cholesky factors of correlation matrices.

When concentration > 1, the distribution favors samples with large diagonal entries (hence large determinant). This is useful when we know a priori that the underlying variables are not correlated.

When concentration < 1, the distribution favors samples with small diagonal entries (hence small determinant). This is useful when we know a priori that some underlying variables are correlated.

### Parameters

- **dimension** (`int`) – dimension of the matrices
- **concentration** (`ndarray`) – concentration/shape parameter of the distribution (often referred to as eta)
- **sample\_method** (`str`) – Either “cvine” or “onion”. Both methods are proposed in [1] and offer the same distribution over correlation matrices. But they are different in how to generate samples. Defaults to “onion”.

### References

[1] *Generating random correlation matrices based on vines and extended onion method*, Daniel Lewandowski, Dorota Kurowicka, Harry Joe

---

```
arg_constraints = {'concentration': <numpyro.distributions.constraints._GreaterThan object>
support = <numpyro.distributions.constraints._CorrCholesky object>
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample\(\)
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log\_prob\(\)
```

## 4.11 LogNormal

```
class LogNormal(loc=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.TransformedDistribution
arg_constraints = {'loc': <numpyro.distributions.constraints._Real object>, 'scale':
reparametrized_params = ['loc', 'scale']
mean
    See numpyro.distributions.distribution.Distribution.mean\(\)
variance
    See numpyro.distributions.distribution.Distribution.variance\(\)
```

## 4.12 Normal

```
class Normal(loc=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'loc': <numpyro.distributions.constraints._Real object>, 'scale':
support = <numpyro.distributions.constraints._Real object>
reparametrized_params = ['loc', 'scale']
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample\(\)
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log\_prob\(\)
mean
    See numpyro.distributions.distribution.Distribution.mean\(\)
variance
    See numpyro.distributions.distribution.Distribution.variance\(\)
```

## 4.13 Pareto

```
class Pareto(alpha, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.TransformedDistribution
arg_constraints = {'alpha': <numpyro.distributions.constraints._GreaterThan object>,
mean
    See numpyro.distributions.distribution.Distribution.mean\(\)
```

**variance**

See `numpyro.distributions.distribution.Distribution.variance()`

**support**

## 4.14 StudentT

```
class StudentT(df=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution

arg_constraints = {'df': <numpyro.distributions.constraints._GreaterThan object>, 'loc': <numpyro.distributions.constraints._Real object>}
support = <numpyro.distributions.constraints._Real object>
reparametrized_params = ['loc', 'scale']
sample(key, sample_shape=())
See numpyro.distributions.distribution.Distribution.sample()

log_prob(value)
See numpyro.distributions.distribution.Distribution.log_prob()

mean
See numpyro.distributions.distribution.Distribution.mean()

variance
See numpyro.distributions.distribution.Distribution.variance()
```

## 4.15 TruncatedCauchy

```
class TruncatedCauchy(low=0.0, loc=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution

arg_constraints = {'loc': <numpyro.distributions.constraints._Real object>, 'low': <numpyro.distributions.constraints._Real object>}
reparametrized_params = ['low', 'loc', 'scale']
sample(key, sample_shape=())
See numpyro.distributions.distribution.Distribution.sample()

log_prob(value)
See numpyro.distributions.distribution.Distribution.log_prob()

mean
See numpyro.distributions.distribution.Distribution.mean()

variance
See numpyro.distributions.distribution.Distribution.variance()

support
```

## 4.16 TruncatedNormal

```
class TruncatedNormal(low=0.0, loc=0.0, scale=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution

arg_constraints = {'loc': <numpyro.distributions.constraints._Real object>, 'low': <numpyro.distributions.constraints._Real object>}
```

```
reparametrized_params = ['low', 'loc', 'scale']
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample\(\)
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log\_prob\(\)
mean
    See numpyro.distributions.distribution.Distribution.mean\(\)
variance
    See numpyro.distributions.distribution.Distribution.variance\(\)
support
```

## 4.17 Uniform

```
class Uniform(low=0.0, high=1.0, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'high': <numpyro.distributions.constraints.\_Dependent object>, 'low': <numpyro.distributions.constraints.\_Dependent object>}
reparametrized_params = ['low', 'high']
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample\(\)
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log\_prob\(\)
mean
    See numpyro.distributions.distribution.Distribution.mean\(\)
variance
    See numpyro.distributions.distribution.Distribution.variance\(\)
support
```



# CHAPTER 5

---

## Discrete Distributions

---

### 5.1 Bernoulli

`Bernoulli` (*probs=None*, *logits=None*, *validate\_args=None*)

### 5.2 BernoulliLogits

```
class BernoulliLogits(logits=None, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'logits': <numpyro.distributions.constraints._Real object>}
    support = <numpyro.distributions.constraints._Boolean object>
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log_prob()
    probs
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

### 5.3 BernoulliProbs

```
class BernoulliProbs(probs, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
```

---

```
arg_constraints = {'probs': <numpyro.distributions.constraints._Interval object>}
support = <numpyro.distributions.constraints._Boolean object>
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample()
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log_prob()
mean
    See numpyro.distributions.distribution.Distribution.mean()
variance
    See numpyro.distributions.distribution.Distribution.variance()
```

## 5.4 Binomial

```
Binomial(total_count=1, probs=None, logits=None, validate_args=None)
```

## 5.5 BinomialLogits

```
class BinomialLogits(logits, total_count=1, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'logits': <numpyro.distributions.constraints._Real object>, 'total_count': <numpyro.distributions.constraints._Interval object>}
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample()
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log_prob()
probs
mean
    See numpyro.distributions.distribution.Distribution.mean()
variance
    See numpyro.distributions.distribution.Distribution.variance()
support
```

## 5.6 BinomialProbs

```
class BinomialProbs(probs, total_count=1, validate_args=None)
Bases: numpyro.distributions.distribution.Distribution
arg_constraints = {'probs': <numpyro.distributions.constraints._Interval object>, 'total_count': <numpyro.distributions.constraints._Interval object>}
sample(key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample()
log_prob(value)
    See numpyro.distributions.distribution.Distribution.log_prob()
```

---

**mean**  
See `numpyro.distributions.distribution.Distribution.mean()`

**variance**  
See `numpyro.distributions.distribution.Distribution.variance()`

**support**

## 5.7 Categorical

`Categorical(probs=None, logits=None, validate_args=None)`

## 5.8 CategoricalLogits

**class CategoricalLogits(logits, validate\_args=None)**  
Bases: `numpyro.distributions.distribution.Distribution`

**arg\_constraints** = {'logits': <`numpyro.distributions.constraints._Real object`>}

**sample** (key, sample\_shape=())  
See `numpyro.distributions.distribution.Distribution.sample()`

**log\_prob** (value)  
See `numpyro.distributions.distribution.Distribution.log_prob()`

**probs**

**mean**  
See `numpyro.distributions.distribution.Distribution.mean()`

**variance**  
See `numpyro.distributions.distribution.Distribution.variance()`

**support**

## 5.9 CategoricalProbs

**class CategoricalProbs(probs, validate\_args=None)**  
Bases: `numpyro.distributions.distribution.Distribution`

**arg\_constraints** = {'probs': <`numpyro.distributions.constraints._Simplex object`>}

**sample** (key, sample\_shape=())  
See `numpyro.distributions.distribution.Distribution.sample()`

**log\_prob** (value)  
See `numpyro.distributions.distribution.Distribution.log_prob()`

**mean**  
See `numpyro.distributions.distribution.Distribution.mean()`

**variance**  
See `numpyro.distributions.distribution.Distribution.variance()`

**support**

## 5.10 Multinomial

```
Multinomial(total_count=1, probs=None, logits=None, validate_args=None)
```

## 5.11 MultinomialLogits

```
class MultinomialLogits(logits, total_count=1, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'logits': <numpyro.distributions.constraints.\_Real object>, 'total_count': <numpyro.distributions.constraints.\_NonNegativeInteger object>}
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample\(\)
    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log\_prob\(\)
    probs
    mean
        See numpyro.distributions.distribution.Distribution.mean\(\)
    variance
        See numpyro.distributions.distribution.Distribution.variance\(\)
    support
```

## 5.12 MultinomialProbs

```
class MultinomialProbs(probs, total_count=1, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'probs': <numpyro.distributions.constraints.\_Simplex object>, 'total_count': <numpyro.distributions.constraints.\_NonNegativeInteger object>}
    sample(key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample\(\)
    log_prob(value)
        See numpyro.distributions.distribution.Distribution.log\_prob\(\)
    mean
        See numpyro.distributions.distribution.Distribution.mean\(\)
    variance
        See numpyro.distributions.distribution.Distribution.variance\(\)
    support
```

## 5.13 Poisson

```
class Poisson(rate, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'rate': <numpyro.distributions.constraints.\_GreaterThan object>}
    support = <numpyro.distributions.constraints.\_IntegerGreaterThan object>
```

**sample** (*key, sample\_shape=()*)  
See `numpyro.distributions.distribution.Distribution.sample()`

**log\_prob** (*value*)  
See `numpyro.distributions.distribution.Distribution.log_prob()`

**mean**  
See `numpyro.distributions.distribution.Distribution.mean()`

**variance**  
See `numpyro.distributions.distribution.Distribution.variance()`



# CHAPTER 6

---

## Pyro Primitives

---

### `param`(*name*, *init\_value*)

Annotate the given site as an optimizable parameter for use with `jax.experimental.optimizers`. For an example of how `param` statements can be used in inference algorithms, refer to `svi()`.

#### Parameters

- `name` (`str`) – name of site.
- `init_value` (`numpy.ndarray`) – initial value specified by the user. Note that the onus of using this to initialize the optimizer is on the user / inference algorithm, since there is no global parameter store in NumPyro.

**Returns** value for the parameter. Unless wrapped inside a handler like `substitute`, this will simply return the initial value.

### `sample`(*name*, *fn*, *obs=None*)

Returns a random sample from the stochastic function *fn*. This can have additional side effects when wrapped inside effect handlers like `substitute`.

#### Parameters

- `name` (`str`) – name of the sample site
- `fn` – Python callable
- `obs` (`numpy.ndarray`) – observed value

**Returns** sample from the stochastic *fn*.



# CHAPTER 7

## Effect Handlers

This provides a small set of effect handlers in NumPyro that are modeled after Pyro's `poutine` module. For a tutorial on effect handlers more generally, readers are encouraged to read [Poutine: A Guide to Programming with Effect Handlers in Pyro](#). These simple effect handlers can be composed together or new ones added to enable implementation of custom inference utilities and algorithms.

### Example

As an example, we are using `seed`, `trace` and `substitute` handlers to define the `log_likelihood` function below. We first create a logistic regression model and sample from the posterior distribution over the regression parameters using `mcmc()`. The `log_likelihood` function uses effect handlers to run the model by substituting sample sites with values from the posterior distribution and computes the log density for a single data point. The `expected_log_likelihood` function computes the log likelihood for each draw from the joint posterior and aggregates the results, but does so by using JAX's auto-vectorize transform called `vmap` so that we do not need to loop over all the data points.

```
>>> N, D = 3000, 3
>>> def logistic_regression(data, labels):
...     coefs = sample('coefs', dist.Normal(np.zeros(D), np.ones(D)))
...     intercept = sample('intercept', dist.Normal(0., 10.))
...     logits = np.sum(coefs * data + intercept, axis=-1)
...     return sample('obs', dist.Bernoulli(logits=logits), obs=labels)

>>> data = random.normal(random.PRNGKey(0), (N, D))
>>> true_coefs = np.arange(1., D + 1.)
>>> logits = np.sum(true_coefs * data, axis=-1)
>>> labels = dist.Bernoulli(logits=logits).sample(random.PRNGKey(1))

>>> init_params, potential_fn, constrain_fn = initialize_model(random.PRNGKey(2),
...     logistic_regression, data, labels)
>>> num_warmup, num_samples = 1000, 1000
>>> samples = mcmc(num_warmup, num_samples, init_params,
...     potential_fn=potential_fn,
...     constrain_fn=constrain_fn)
warmup: 100%|| 1000/1000 [00:09<00:00, 109.40it/s, 1 steps of size 5.83e-01. acc.=
prob=0.79]
```

(continues on next page)

(continued from previous page)

```

sample: 100%|| 1000/1000 [00:00<00:00, 1252.39it/s, 1 steps of size 5.83e-01. acc. ↵
prob=0.85]

          mean      sd    5.5%   94.5%    n_eff    Rhat
coefs[0]    0.96    0.07    0.85    1.07    455.35    1.01
coefs[1]    2.05    0.09    1.91    2.20    332.00    1.01
coefs[2]    3.18    0.13    2.96    3.37    320.27    1.00
intercept   -0.03    0.02   -0.06    0.00    402.53    1.00

>>> def log_likelihood(rng, params, model, *args, **kwargs):
...     model = substitute(seed(model, rng), params)
...     model_trace = trace(model).get_trace(*args, **kwargs)
...     obs_node = model_trace['obs']
...     return np.sum(obs_node['fn'].log_prob(obs_node['value']))

>>> def expected_log_likelihood(rng, params, model, *args, **kwargs):
...     n = list(params.values())[0].shape[0]
...     log_lk_fn = vmap(lambda rng, params: log_likelihood(rng, params, model, *args,
...     **kwargs))
...     log_lk_vals = log_lk_fn(random.split(rng, n), params)
...     return logsumexp(log_lk_vals) - np.log(n)

>>> print(expected_log_likelihood(random.PRNGKey(2), samples, logistic_regression,
... data, labels))
-876.172

```

**class** `block`(*fn=None, hide\_fn=<function block.<lambda>>*)  
Bases: `numpyro.handlers.Messenger`

Given a callable *fn*, return another callable that selectively hides primitive sites where *hide\_fn* returns True from other effect handlers on the stack.

#### Parameters

- **fn** – Python callable with NumPyro primitives.
- **hide\_fn** – function which when given a dictionary containing site-level metadata returns whether it should be blocked.

#### Example:

```

>>> def model():
...     a = sample('a', dist.Normal(0., 1.))
...     return sample('b', dist.Normal(a, 1.))

>>> model = seed(model, random.PRNGKey(0))
>>> block_all = block(model)
>>> block_a = block(model, lambda site: site['name'] == 'a')
>>> trace_block_all = trace(block_all).get_trace()
>>> assert not {'a', 'b'}.intersection(trace_block_all.keys())
>>> trace_block_a = trace(block_a).get_trace()
>>> assert 'a' not in trace_block_a
>>> assert 'b' in trace_block_a

```

`process_message(msg)`

**class** `replay`(*fn, guide\_trace*)  
Bases: `numpyro.handlers.Messenger`

Given a callable *fn* and an execution trace *guide\_trace*, return a callable which substitutes *sample* calls in *fn* with values from the corresponding site names in *guide\_trace*.

### Parameters

- **fn** – Python callable with NumPyro primitives.
- **guide\_trace** – an OrderedDict containing execution metadata.

### Example

```
>>> def model():
...     sample('a', dist.Normal(0., 1.))

>>> exec_trace = trace(seed(model, random.PRNGKey(0))).get_trace()
>>> print(exec_trace['a']['value'])
-0.20584235
>>> replayed_trace = trace(replay(model, exec_trace)).get_trace()
>>> print(exec_trace['a']['value'])
-0.20584235
>>> assert replayed_trace['a']['value'] == exec_trace['a']['value']
```

**process\_message** (*msg*)

**class seed** (*fn, rng*)

Bases: numpyro.handlers.Messenger

JAX uses a functional pseudo random number generator that requires passing in a seed `PRNGKey()` to every stochastic function. The `seed` handler allows us to initially seed a stochastic function with a `PRNGKey()`. Every call to the `sample()` primitive inside the function results in a splitting of this initial seed so that we use a fresh seed for each subsequent call without having to explicitly pass in a `PRNGKey` to each `sample` call.

**process\_message** (*msg*)

**class substitute** (*fn=None, param\_map=None*)

Bases: numpyro.handlers.Messenger

Given a callable *fn* and a dict *param\_map* keyed by site names, return a callable which substitutes all primitive calls in *fn* with values from *param\_map* whose key matches the site name. If the site name is not present in *param\_map*, there is no side effect.

### Parameters

- **fn** – Python callable with NumPyro primitives.
- **param\_map** (*dict*) – dictionary of `numpy.ndarray` values keyed by site names.

### Example:

```
>>> def model():
...     sample('a', dist.Normal(0., 1.))

>>> model = seed(model, random.PRNGKey(0))
>>> exec_trace = trace(substitute(model, {'a': -1})).get_trace()
>>> assert exec_trace['a']['value'] == -1
```

**process\_message** (*msg*)

**class trace** (*fn=None*)

Bases: numpyro.handlers.Messenger

Returns a handler that records the inputs and outputs at primitive calls inside *fn*.

## Example

```
>>> def model():
...     sample('a', dist.Normal(0., 1.))

>>> exec_trace = trace(seed(model, random.PRNGKey(0))).get_trace()
>>> pp pprint(exec_trace)
OrderedDict([('a',
    {'args': (),
     'fn': <numpyro.distributions.continuous.Normal object at 0x7f9e689b1eb8>,
     'is_observed': False,
     'kwargs': {'random_state': DeviceArray([0, 0], dtype=uint32)},
     'name': 'a',
     'type': 'sample',
     'value': DeviceArray(-0.20584235, dtype=float32)}])
```

**postprocess\_message** (msg)

**get\_trace** (\*args, \*\*kwargs)

Run the wrapped callable and return the recorded trace.

### Parameters

- **\*args** – arguments to the callable.
- **\*\*kwargs** – keyword arguments to the callable.

**Returns** *OrderedDict* containing the execution trace.

# CHAPTER 8

---

## Autocorrelation

---

**autocorrelation**(*x*, *axis*=0)

Computes the autocorrelation of samples at dimension *axis*.

### Parameters

- **x** (`numpy.ndarray`) – the input array.
- **axis** (`int`) – the dimension to calculate autocorrelation.

**Returns** autocorrelation of *x*.

**Return type** `numpy.ndarray`



# CHAPTER 9

---

## Autocovariance

---

**autocovariance**(*x*, *axis*=0)

Computes the autocovariance of samples at dimension *axis*.

### Parameters

- **x** (`numpy.ndarray`) – the input array.
- **axis** (`int`) – the dimension to calculate autocovariance.

**Returns** autocovariance of *x*.

**Return type** `numpy.ndarray`



# CHAPTER 10

---

## Effective Sample Size

---

`effective_sample_size(x)`

Computes effective sample size of input `x`, where the first dimension of `x` is chain dimension and the second dimension of `x` is draw dimension.

**References:**

1. *Introduction to Markov Chain Monte Carlo*, Charles J. Geyer
2. *Stan Reference Manual version 2.18*, Stan Development Team

**Parameters** `x` (`numpy.ndarray`) – the input array.

**Returns** effective sample size of `x`.

**Return type** `numpy.ndarray`



# CHAPTER 11

---

## Gelman Rubin

---

**gelman\_rubin**(*x*)

Computes R-hat over chains of samples *x*, where the first dimension of *x* is chain dimension and the second dimension of *x* is draw dimension. It is required that `input.shape[0] >= 2` and `input.shape[1] >= 2`.

**Parameters** **x** (`numpy.ndarray`) – the input array.

**Returns** R-hat of *x*.

**Return type** `numpy.ndarray`



# CHAPTER 12

---

## Split Gelman Rubin

---

`split_gelman_rubin(x)`

Computes split R-hat over chains of samples `x`, where the first dimension of `x` is chain dimension and the second dimension of `x` is draw dimension. It is required that `input.shape[1] >= 4`.

**Parameters** `x` (`numpy.ndarray`) – the input array.

**Returns** split R-hat of `x`.

**Return type** `numpy.ndarray`



# CHAPTER 13

---

## HPDI

---

**hpdi** (*x, prob=0.89, axis=0*)

Computes “highest posterior density interval” (HPDI) which is the narrowest interval with probability mass *prob*.

### Parameters

- **x** (`numpy.ndarray`) – the input array.
- **prob** (`float`) – the probability mass of samples within the interval.
- **axis** (`int`) – the dimension to calculate hpdi.

**Returns** quantiles of input at  $(1 - \text{probs}) / 2$  and  $(1 + \text{probs}) / 2$ .

**Return type** `numpy.ndarray`



# CHAPTER 14

---

## Summary

---

**summary** (*samples*, *prob*=0.89)

Prints a summary table displaying diagnostics of *samples* from the posterior. The diagnostics displayed are mean, standard deviation, the 89% Credibility Interval, *effective\_sample\_size()*, *split\_gelman\_rubin()*.

### Parameters

- **samples** – a collection of input samples.
- **prob** (*float*) – the probability mass of samples within the HPDI interval.



# CHAPTER 15

---

## Indices and tables

---

- genindex
- search



---

## Python Module Index

---

n

`numpyro.handlers`, 29



---

## Index

---

### A

arg\_constraints (*BernoulliLogits attribute*), 21  
arg\_constraints (*BernoulliProbs attribute*), 21  
arg\_constraints (*Beta attribute*), 13  
arg\_constraints (*BinomialLogits attribute*), 22  
arg\_constraints (*BinomialProbs attribute*), 22  
arg\_constraints (*CategoricalLogits attribute*), 23  
arg\_constraints (*CategoricalProbs attribute*), 23  
arg\_constraints (*Cauchy attribute*), 13  
arg\_constraints (*Chi2 attribute*), 14  
arg\_constraints (*Dirichlet attribute*), 14  
arg\_constraints (*Distribution attribute*), 9  
arg\_constraints (*Exponential attribute*), 14  
arg\_constraints (*Gamma attribute*), 15  
arg\_constraints (*GaussianRandomWalk attribute*), 15  
arg\_constraints (*HalfCauchy attribute*), 15  
arg\_constraints (*HalfNormal attribute*), 16  
arg\_constraints (*LKJCholesky attribute*), 16  
arg\_constraints (*LogNormal attribute*), 17  
arg\_constraints (*MultinomialLogits attribute*), 24  
arg\_constraints (*MultinomialProbs attribute*), 24  
arg\_constraints (*Normal attribute*), 17  
arg\_constraints (*Pareto attribute*), 17  
arg\_constraints (*Poisson attribute*), 24  
arg\_constraints (*StudentT attribute*), 18  
arg\_constraints (*TransformedDistribution attribute*), 10  
arg\_constraints (*TruncatedCauchy attribute*), 18  
arg\_constraints (*TruncatedNormal attribute*), 18  
arg\_constraints (*Uniform attribute*), 19  
autocorrelation() (in module *numpyro.diagnostics*), 33  
autocovariance() (in module *numpyro.diagnostics*), 35

### B

batch\_shape (*Distribution attribute*), 9

Bernoulli() (in module *numpyro.distributions.discrete*), 21  
BernoulliLogits (class in *numpyro.distributions.discrete*), 21  
BernoulliProbs (class in *numpyro.distributions.discrete*), 21  
Beta (class in *numpyro.distributions.continuous*), 13  
Binomial() (in module *numpyro.distributions.discrete*), 22

BinomialLogits (class in *numpyro.distributions.discrete*), 22  
BinomialProbs (class in *numpyro.distributions.discrete*), 22  
block (class in *numpyro.handlers*), 30

### C

Categorical() (in module *numpyro.distributions.discrete*), 23  
CategoricalLogits (class in *numpyro.distributions.discrete*), 23  
CategoricalProbs (class in *numpyro.distributions.discrete*), 23  
Cauchy (class in *numpyro.distributions.continuous*), 13  
Chi2 (class in *numpyro.distributions.continuous*), 14

### D

Dirichlet (class in *numpyro.distributions.continuous*), 14  
Distribution (class in *numpyro.distributions.distribution*), 9

### E

effective\_sample\_size() (in module *numpyro.diagnostics*), 37  
elbo() (in module *numpyro.svi*), 8  
evaluate() (in module *numpyro.svi.svi*), 8  
event\_shape (*Distribution attribute*), 10  
Exponential (class in *numpyro.distributions.continuous*), 14

## F

`fori_collect()` (*in module numpyro.util*), 5

## G

`Gamma` (*class in numpyro.distributions.continuous*), 15  
`GaussianRandomWalk` (*class in numpyro.distributions.continuous*), 15  
`gelman_rubin()` (*in module numpyro.diagnostics*), 39  
`get_trace()` (*trace method*), 32

## H

`HalfCauchy` (*class in numpyro.distributions.continuous*), 15  
`HalfNormal` (*class in numpyro.distributions.continuous*), 16  
`hmc()` (*in module numpyro.mcmc*), 2  
`HMCState` (*in module numpyro.mcmc*), 4  
`hpdi()` (*in module numpyro.diagnostics*), 43

## I

`init_fn()` (*in module numpyro.svi.svi*), 7  
`init_kernel()` (*in module numpyro.mcmc.hmc*), 3  
`initialize_model()` (*in module numpyro.hmc\_util*), 4  
`is_reparametrized` (*TransformedDistribution attribute*), 10

## L

`LKJCholesky` (*class in numpyro.distributions.continuous*), 16  
`log_prob()` (*BernoulliLogits method*), 21  
`log_prob()` (*BernoulliProbs method*), 22  
`log_prob()` (*Beta method*), 13  
`log_prob()` (*BinomialLogits method*), 22  
`log_prob()` (*BinomialProbs method*), 22  
`log_prob()` (*CategoricalLogits method*), 23  
`log_prob()` (*CategoricalProbs method*), 23  
`log_prob()` (*Cauchy method*), 13  
`log_prob()` (*Dirichlet method*), 14  
`log_prob()` (*Distribution method*), 10  
`log_prob()` (*Exponential method*), 14  
`log_prob()` (*Gamma method*), 15  
`log_prob()` (*GaussianRandomWalk method*), 15  
`log_prob()` (*HalfCauchy method*), 15  
`log_prob()` (*HalfNormal method*), 16  
`log_prob()` (*LKJCholesky method*), 17  
`log_prob()` (*MultinomialLogits method*), 24  
`log_prob()` (*MultinomialProbs method*), 24  
`log_prob()` (*Normal method*), 17  
`log_prob()` (*Poisson method*), 25  
`log_prob()` (*StudentT method*), 18  
`log_prob()` (*TransformedDistribution method*), 11

`log_prob()` (*TruncatedCauchy method*), 18  
`log_prob()` (*TruncatedNormal method*), 19  
`log_prob()` (*Uniform method*), 19  
`LogNormal` (*class in numpyro.distributions.continuous*), 17

## M

`mcmc()` (*in module numpyro.mcmc*), 1  
`mean` (*BernoulliLogits attribute*), 21  
`mean` (*BernoulliProbs attribute*), 22  
`mean` (*Beta attribute*), 13  
`mean` (*BinomialLogits attribute*), 22  
`mean` (*BinomialProbs attribute*), 22  
`mean` (*CategoricalLogits attribute*), 23  
`mean` (*CategoricalProbs attribute*), 23  
`mean` (*Cauchy attribute*), 14  
`mean` (*Dirichlet attribute*), 14  
`mean` (*Distribution attribute*), 10  
`mean` (*Exponential attribute*), 14  
`mean` (*Gamma attribute*), 15  
`mean` (*GaussianRandomWalk attribute*), 15  
`mean` (*HalfCauchy attribute*), 16  
`mean` (*HalfNormal attribute*), 16  
`mean` (*LogNormal attribute*), 17  
`mean` (*MultinomialLogits attribute*), 24  
`mean` (*MultinomialProbs attribute*), 24  
`mean` (*Normal attribute*), 17  
`mean` (*Pareto attribute*), 17  
`mean` (*Poisson attribute*), 25  
`mean` (*StudentT attribute*), 18  
`mean` (*TransformedDistribution attribute*), 11  
`mean` (*TruncatedCauchy attribute*), 18  
`mean` (*TruncatedNormal attribute*), 19  
`mean` (*Uniform attribute*), 19  
`Multinomial` (*class in module numpyro.distributions.discrete*), 24  
`MultinomialLogits` (*class in numpyro.distributions.discrete*), 24  
`MultinomialProbs` (*class in numpyro.distributions.discrete*), 24

## N

`Normal` (*class in numpyro.distributions.continuous*), 17  
`numpyro.handlers` (*module*), 29

## P

`param()` (*in module numpyro.handlers*), 27  
`Pareto` (*class in numpyro.distributions.continuous*), 17  
`Poisson` (*class in numpyro.distributions.discrete*), 24  
`postprocess_message()` (*trace method*), 32  
`probs` (*BernoulliLogits attribute*), 21  
`probs` (*BinomialLogits attribute*), 22  
`probs` (*CategoricalLogits attribute*), 23  
`probs` (*MultinomialLogits attribute*), 24

`process_message ()` (*block method*), 30  
`process_message ()` (*replay method*), 31  
`process_message ()` (*seed method*), 31  
`process_message ()` (*substitute method*), 31

**R**

`reparametrized_params` (*Cauchy attribute*), 13  
`reparametrized_params` (*Distribution attribute*), 9  
`reparametrized_params` (*Exponential attribute*), 14  
`reparametrized_params` (*GaussianRandomWalk attribute*), 15  
`reparametrized_params` (*HalfCauchy attribute*), 15  
`reparametrized_params` (*HalfNormal attribute*), 16  
`reparametrized_params` (*LogNormal attribute*), 17  
`reparametrized_params` (*Normal attribute*), 17  
`reparametrized_params` (*StudentT attribute*), 18  
`reparametrized_params` (*TruncatedCauchy attribute*), 18  
`reparametrized_params` (*TruncatedNormal attribute*), 18  
`reparametrized_params` (*Uniform attribute*), 19  
`replay` (*class in numpyro.handlers*), 30

**S**

`sample ()` (*BernoulliLogits method*), 21  
`sample ()` (*BernoulliProbs method*), 22  
`sample ()` (*Beta method*), 13  
`sample ()` (*BinomialLogits method*), 22  
`sample ()` (*BinomialProbs method*), 22  
`sample ()` (*CategoricalLogits method*), 23  
`sample ()` (*CategoricalProbs method*), 23  
`sample ()` (*Cauchy method*), 13  
`sample ()` (*Dirichlet method*), 14  
`sample ()` (*Distribution method*), 10  
`sample ()` (*Exponential method*), 14  
`sample ()` (*Gamma method*), 15  
`sample ()` (*GaussianRandomWalk method*), 15  
`sample ()` (*in module numpyro.handlers*), 27  
`sample ()` (*LKJCholesky method*), 17  
`sample ()` (*MultinomialLogits method*), 24  
`sample ()` (*MultinomialProbs method*), 24  
`sample ()` (*Normal method*), 17  
`sample ()` (*Poisson method*), 24  
`sample ()` (*StudentT method*), 18  
`sample ()` (*TransformedDistribution method*), 10  
`sample ()` (*TruncatedCauchy method*), 18  
`sample ()` (*TruncatedNormal method*), 19  
`sample ()` (*Uniform method*), 19  
`sample_kernel ()` (*in module numpyro.mcmc.hmc*), 4  
`seed` (*class in numpyro.handlers*), 31

`split_gelman_rubin ()` (*in module numpyro.diagnostics*), 41  
`StudentT` (*class in numpyro.distributions.continuous*), 18  
`substitute` (*class in numpyro.handlers*), 31  
`summary ()` (*in module numpyro.diagnostics*), 45  
`support` (*BernoulliLogits attribute*), 21  
`support` (*BernoulliProbs attribute*), 22  
`support` (*Beta attribute*), 13  
`support` (*BinomialLogits attribute*), 22  
`support` (*BinomialProbs attribute*), 23  
`support` (*CategoricalLogits attribute*), 23  
`support` (*CategoricalProbs attribute*), 23  
`support` (*Cauchy attribute*), 13  
`support` (*Dirichlet attribute*), 14  
`support` (*Distribution attribute*), 9  
`support` (*Exponential attribute*), 14  
`support` (*Gamma attribute*), 15  
`support` (*GaussianRandomWalk attribute*), 15  
`support` (*LKJCholesky attribute*), 17  
`support` (*MultinomialLogits attribute*), 24  
`support` (*MultinomialProbs attribute*), 24  
`support` (*Normal attribute*), 17  
`support` (*Pareto attribute*), 18  
`support` (*Poisson attribute*), 24  
`support` (*StudentT attribute*), 18  
`support` (*TransformedDistribution attribute*), 10  
`support` (*TruncatedCauchy attribute*), 18  
`support` (*TruncatedNormal attribute*), 19  
`support` (*Uniform attribute*), 19  
`svi ()` (*in module numpyro.svi*), 7

**T**

`trace` (*class in numpyro.handlers*), 31  
`TransformedDistribution` (*class in numpyro.distributions.distribution*), 10  
`TruncatedCauchy` (*class in numpyro.distributions.continuous*), 18  
`TruncatedNormal` (*class in numpyro.distributions.continuous*), 18

**U**

`Uniform` (*class in numpyro.distributions.continuous*), 19  
`update_fn ()` (*in module numpyro.svi.svi*), 7

**V**

`variance` (*BernoulliLogits attribute*), 21  
`variance` (*BernoulliProbs attribute*), 22  
`variance` (*Beta attribute*), 13  
`variance` (*BinomialLogits attribute*), 22  
`variance` (*BinomialProbs attribute*), 23  
`variance` (*CategoricalLogits attribute*), 23  
`variance` (*CategoricalProbs attribute*), 23

variance (*Cauchy attribute*), 14  
variance (*Dirichlet attribute*), 14  
variance (*Distribution attribute*), 10  
variance (*Exponential attribute*), 14  
variance (*Gamma attribute*), 15  
variance (*GaussianRandomWalk attribute*), 15  
variance (*HalfCauchy attribute*), 16  
variance (*HalfNormal attribute*), 16  
variance (*LogNormal attribute*), 17  
variance (*MultinomialLogits attribute*), 24  
variance (*MultinomialProbs attribute*), 24  
variance (*Normal attribute*), 17  
variance (*Pareto attribute*), 17  
variance (*Poisson attribute*), 25  
variance (*StudentT attribute*), 18  
variance (*TransformedDistribution attribute*), 11  
variance (*TruncatedCauchy attribute*), 18  
variance (*TruncatedNormal attribute*), 19  
variance (*Uniform attribute*), 19