
NumPyro Documentation

Release 0.0

Uber AI Labs

Nov 01, 2019

1	Pyro Primitives	1
1.1	param	1
1.2	sample	1
1.3	plate	2
1.4	module	2
2	Effect Handlers	3
2.1	block	4
2.2	condition	5
2.3	replay	5
2.4	scale	6
2.5	seed	6
2.6	substitute	6
2.7	trace	7
3	Base Distribution	9
3.1	Distribution	9
3.2	Independent	11
3.3	TransformedDistribution	11
4	Continuous Distributions	13
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	InverseGamma	16
4.11	LKJ	17
4.12	LKJCholesky	17
4.13	LogNormal	18
4.14	MultivariateNormal	18
4.15	LowRankMultivariateNormal	19
4.16	Normal	19

4.17	Pareto	20
4.18	StudentT	20
4.19	TruncatedCauchy	20
4.20	TruncatedNormal	21
4.21	Uniform	21
5	Discrete Distributions	23
5.1	Bernoulli	23
5.2	BernoulliLogits	23
5.3	BernoulliProbs	23
5.4	BetaBinomial	24
5.5	Binomial	24
5.6	BinomialLogits	24
5.7	BinomialProbs	25
5.8	Categorical	25
5.9	CategoricalLogits	25
5.10	CategoricalProbs	26
5.11	Delta	26
5.12	GammaPoisson	26
5.13	Multinomial	27
5.14	MultinomialLogits	27
5.15	MultinomialProbs	27
5.16	OrderedLogistic	28
5.17	Poisson	28
5.18	PRNGIdentity	28
5.19	ZeroInflatedPoisson	29
6	Constraints	31
6.1	boolean	31
6.2	corr_cholesky	31
6.3	corr_matrix	31
6.4	dependent	31
6.5	greater_than	31
6.6	integer_interval	31
6.7	integer_greater_than	32
6.8	interval	32
6.9	lower_cholesky	32
6.10	multinomial	32
6.11	nonnegative_integer	32
6.12	ordered_vector	32
6.13	positive	32
6.14	positive_definite	32
6.15	positive_integer	32
6.16	real	32
6.17	real_vector	33
6.18	simplex	33
6.19	unit_interval	33
7	Transforms	35
7.1	biject_to	35
7.2	Transform	35
7.3	AbsTransform	35
7.4	AffineTransform	36
7.5	ComposeTransform	36

7.6	CorrCholeskyTransform	36
7.7	ExpTransform	37
7.8	IdentityTransform	37
7.9	InvCholeskyTransform	37
7.10	LowerCholeskyTransform	37
7.11	MultivariateAffineTransform	38
7.12	OrderedTransform	38
7.13	PermuteTransform	38
7.14	PowerTransform	39
7.15	SigmoidTransform	39
7.16	StickBreakingTransform	39
8	Flows	41
8.1	InverseAutoregressiveTransform	41
9	Markov Chain Monte Carlo (MCMC)	43
9.1	Hamiltonian Monte Carlo	43
9.2	MCMC Utilities	48
10	Stochastic Variational Inference (SVI)	51
10.1	ELBO	52
10.2	RenyiELBO	53
11	Automatic Guide Generation	55
11.1	AutoContinuous	55
11.2	AutoDiagonalNormal	56
11.3	AutoMultivariateNormal	57
11.4	AutoIAFNormal	57
11.5	AutoLaplaceApproximation	58
11.6	AutoContinuousELBO	58
12	Optimizers	59
12.1	Adam	59
12.2	Adagrad	60
12.3	ClippedAdam	60
12.4	Momentum	61
12.5	RMSPProp	61
12.6	RMSPPropMomentum	62
12.7	SGD	62
12.8	SM3	63
13	Diagnostics	65
13.1	Autocorrelation	65
13.2	Autocovariance	65
13.3	Effective Sample Size	66
13.4	Gelman Rubin	66
13.5	Split Gelman Rubin	66
13.6	HPDI	66
13.7	Summary	67
14	Runtime Utilities	69
14.1	enable_validation	69
14.2	validation_enabled	69
14.3	set_platform	69
14.4	set_host_device_count	70

15 Inference Utilities	71
15.1 Predictive	71
15.2 log_density	72
15.3 transform_fn	72
15.4 constrain_fn	72
15.5 potential_energy	73
15.6 log_likelihood	73
15.7 find_valid_initial_params	74
15.8 Initialization Strategies	74
16 Indices and tables	77
Python Module Index	79
Index	81

1.1 param

param (*name*, *init_value=None*, ***kwargs*)

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.

1.2 sample

sample (*name*, *fn*, *obs=None*, *rng_key=None*, *sample_shape=()*)

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

Note: By design, *sample* primitive is meant to be used inside a NumPyro model. Then *seed* handler is used to inject a random state to *fn*. In those situations, *rng_key* keyword will take no effect.

Parameters

- **name** (*str*) – name of the sample site
- **fn** – Python callable

- **obs** (*numpy.ndarray*) – observed value
- **rng_key** (*jax.random.PRNGKey*) – an optional random key for *fn*.
- **sample_shape** – Shape of samples to be drawn.

Returns sample from the stochastic *fn*.

1.3 plate

class plate (*name, size, subsample_size=None, dim=None*)

Construct for annotating conditionally independent variables. Within a *plate* context manager, *sample* sites will be automatically broadcasted to the size of the plate. Additionally, a scale factor might be applied by certain inference algorithms if *subsample_size* is specified.

Parameters

- **name** (*str*) – Name of the plate.
- **size** (*int*) – Size of the plate.
- **subsample_size** (*int*) – Optional argument denoting the size of the mini-batch. This can be used to apply a scaling factor by inference algorithms. e.g. when computing ELBO using a mini-batch.
- **dim** (*int*) – Optional argument to specify which dimension in the tensor is used as the plate dim. If *None* (default), the leftmost available dim is allocated.

1.4 module

module (*name, nn, input_shape=None*)

Declare a *stax* style neural network inside a model so that its parameters are registered for optimization via *param()* statements.

Parameters

- **name** (*str*) – name of the module to be registered.
- **nn** (*tuple*) – a tuple of (*init_fn, apply_fn*) obtained by a *stax* constructor function.
- **input_shape** (*tuple*) – shape of the input taken by the neural network.

Returns a *apply_fn* with bound parameters that takes an array as an input and returns the neural network transformed output array.

CHAPTER 2

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 `log_predictive_density` function computes the log likelihood for each draw from the joint posterior and aggregates the results for all the data points, 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 = numpyro.sample('coefs', dist.Normal(np.zeros(D), np.ones(D)))
...     intercept = numpyro.sample('intercept', dist.Normal(0., 10.))
...     logits = np.sum(coefs * data + intercept, axis=-1)
...     return numpyro.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))

>>> num_warmup, num_samples = 1000, 1000
>>> mcmc = MCMC(NUTS(model=logistic_regression), num_warmup, num_samples)
>>> mcmc.run(random.PRNGKey(2), data, labels)
sample: 100%| 1000/1000 [00:00<00:00, 1252.39it/s, 1 steps of size 5.83e-01. acc.
↳prob=0.85]
>>> mcmc.print_summary()
```

(continues on next page)

(continued from previous page)

	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_key, params, model, *args, **kwargs):
...     model = handlers.substitute(handlers.seed(model, rng_key), params)
...     model_trace = handlers.trace(model).get_trace(*args, **kwargs)
...     obs_node = model_trace['obs']
...     return obs_node['fn'].log_prob(obs_node['value'])

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

>>> print(log_predictive_density(random.PRNGKey(2), mcmc.get_samples(),
...     logistic_regression, data, labels))
-874.89813

```

2.1 block

class block (*fn=None, hide_fn=<function block.<lambda>>>*)

Bases: `numpyro.primitives.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 = numpyro.sample('a', dist.Normal(0., 1.))
...     return numpyro.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*)

2.2 condition

class condition (*fn=None, param_map=None, substitute_fn=None*)

Bases: `numpyro.primitives.Messenger`

Conditions unobserved sample sites to values from *param_map* or *condition_fn*. Similar to *substitute* except that it only affects *sample* sites and changes the *is_observed* property to *True*.

Parameters

- **fn** – Python callable with NumPyro primitives.
- **param_map** (*dict*) – dictionary of *numpy.ndarray* values keyed by site names.
- **condition_fn** – callable that takes in a site dict and returns a numpy array or *None* (in which case the handler has no side effect).

Example:

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

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

process_message (*msg*)

2.3 replay

class replay (*fn, guide_trace*)

Bases: `numpyro.primitives.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():
...     numpyro.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*)

2.4 scale

class scale (*fn=None, scale_factor=1.0*)

Bases: `numpyro.primitives.Messenger`

This messenger rescales the log probability score.

This is typically used for data subsampling or for stratified sampling of data (e.g. in fraud detection where negatives vastly outnumber positives).

Parameters `scale_factor` (*float*) – a positive scaling factor

`process_message` (*msg*)

2.5 seed

class seed (*fn=None, rng_seed=None, rng=None*)

Bases: `numpyro.primitives.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.

Parameters

- **fn** – Python callable with NumPyro primitives.
- **rng_seed** (*int, np.ndarray scalar, or jax.random.PRNGKey*) – a random number generator seed.

Note: Unlike in Pyro, `numpyro.sample` primitive cannot be used without wrapping it in seed handler since there is no global random state. As such, users need to use *seed* as a contextmanager to generate samples from distributions or as a decorator for their model callable (See below).

Example:

```
>>> # as context manager
>>> with handlers.seed(rng_seed=1):
...     x = numpyro.sample('x', dist.Normal(0., 1.))

>>> def model():
...     return numpyro.sample('y', dist.Normal(0., 1.))

>>> # as function decorator (/modifier)
>>> y = handlers.seed(model, rng_seed=1)()
>>> assert x == y
```

`process_message` (*msg*)

2.6 substitute

class substitute (*fn=None, param_map=None, base_param_map=None, substitute_fn=None*)

Bases: `numpyro.primitives.Messenger`

Given a callable *fn* and a dict *param_map* keyed by site names (alternatively, a callable *substitute_fn*), 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.

If a *substitute_fn* is provided, then the value at the site is replaced by the value returned from the call to *substitute_fn* for the given site.

Parameters

- **fn** – Python callable with NumPyro primitives.
- **param_map** (*dict*) – dictionary of *numpy.ndarray* values keyed by site names.
- **base_param_map** (*dict*) – similar to *param_map* but only holds samples from base distributions.
- **substitute_fn** – callable that takes in a site dict and returns a numpy array or *None* (in which case the handler has no side effect).

Example:

```
>>> def model():
...     numpyro.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*)

2.7 trace

class trace (*fn=None*)

Bases: *numpyro.primitives.Messenger*

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

Example

```
>>> def model():
...     numpyro.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': {'rng_key': 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.

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 = []
```

```
static set_default_validate_args(value)
```

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* key to be used for the distribution.
- **sample_shape** (*tuple*) – the sample shape for the distribution.

Returns an array of shape *sample_shape* + *batch_shape* + *event_shape*

Return type `numpy.ndarray`

sample_with_intermediates (*key*, *sample_shape*=())

Same as `sample` except that any intermediate computations are returned (useful for *TransformedDistribution*).

Parameters

- **key** (*jax.random.PRNGKey*) – the *rng_key* key to be used for the distribution.
- **sample_shape** (*tuple*) – the sample shape for the distribution.

Returns an array of shape *sample_shape* + *batch_shape* + *event_shape*

Return type `numpy.ndarray`

transform_with_intermediates (*base_value*)**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 an array with shape *value.shape[:-self.event_shape]*

Return type `numpy.ndarray`

mean

Mean of the distribution.

variance

Variance of the distribution.

to_event (*reinterpreted_batch_ndims*=None)

Interpret the rightmost *reinterpreted_batch_ndims* batch dimensions as dependent event dimensions.

Parameters *reinterpreted_batch_ndims* – Number of rightmost batch dims to interpret as event dims.

Returns An instance of *Independent* distribution.

Return type *Independent*

3.2 Independent

class Independent (*base_dist, reinterpreted_batch_ndims, validate_args=None*)

Bases: *numpyro.distributions.distribution.Distribution*

Reinterprets batch dimensions of a distribution as event dims by shifting the batch-event dim boundary further to the left.

From a practical standpoint, this is useful when changing the result of *log_prob()*. For example, a univariate Normal distribution can be interpreted as a multivariate Normal with diagonal covariance:

```
>>> normal = dist.Normal(np.zeros(3), np.ones(3))
>>> [normal.batch_shape, normal.event_shape]
[(3,), ()]
>>> diag_normal = dist.Independent(normal, 1)
>>> [diag_normal.batch_shape, diag_normal.event_shape]
[(), (3,)]
```

Parameters

- **base_distribution** (*numpyro.distribution.Distribution*) – a distribution instance.
- **reinterpreted_batch_ndims** (*int*) – the number of batch dims to reinterpret as event dims.

arg_constraints = {}

support

reparameterized_params

mean

variance

sample (*key, sample_shape=()*)

log_prob (*value*)

3.3 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  
sample (key, sample_shape=())  
    See numpyro.distributions.distribution.Distribution.sample()  
sample_with_intermediates (key, sample_shape=())  
    See numpyro.distributions.distribution.Distribution.sample_with_intermediates()  
transform_with_intermediates (base_value)  
log_prob (*args, **kwargs)  
    See numpyro.distributions.distribution.Distribution.log_prob()  
mean  
variance
```

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 (*args, **kwargs)
        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(*args, **kwargs)
    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(*args, **kwargs)
        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(*args, **kwargs)
        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>`

reparametrized_params = ['rate']

sample (*key, sample_shape=()*)

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

log_prob (**args, **kwargs*)

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._IntegerGreaterThan object>`

support = `<numpyro.distributions.constraints._RealVector object>`

reparametrized_params = ['scale']

sample (*key, sample_shape=()*)

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

log_prob (**args, **kwargs*)

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.Distribution`

reparametrized_params = ['scale']

support = `<numpyro.distributions.constraints._GreaterThan object>`

```
arg_constraints = {'scale': <numpyro.distributions.constraints._GreaterThan object>}
sample (key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample()
log_prob (*args, **kwargs)
    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.Distribution
    reparametrized_params = ['scale']
    support = <numpyro.distributions.constraints._GreaterThan object>
    arg_constraints = {'scale': <numpyro.distributions.constraints._GreaterThan object>}
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

4.10 InverseGamma

```
class InverseGamma (concentration, rate=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.TransformedDistribution
    arg_constraints = {'concentration': <numpyro.distributions.constraints._GreaterThan object>}
    support = <numpyro.distributions.constraints._GreaterThan object>
    reparametrized_params = ['rate']
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

4.11 LKJ

class `LKJ` (*dimension*, *concentration=1.0*, *sample_method='onion'*, *validate_args=None*)

Bases: `numpyro.distributions.distribution.TransformedDistribution`

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

When `concentration > 1`, the distribution favors samples with large large determinent. This is useful when we know a priori that the underlying variables are not correlated.

When `concentration < 1`, the distribution favors samples with 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._CorrMatrix object>`

mean

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

variance

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

4.12 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 η to make the probability of the correlation matrix M generated from a Cholesky factor propotional 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 (*args, **kwargs)
    See numpyro.distributions.distribution.Distribution.log_prob()
```

4.13 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()
```

4.14 MultivariateNormal

```
class MultivariateNormal (loc=0.0, covariance_matrix=None, precision_matrix=None,
                           scale_tril=None, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'covariance_matrix': <numpyro.distributions.constraints._PositiveDefinite object>,
    support = <numpyro.distributions.constraints._RealVector object>
    reparametrized_params = ['loc', 'covariance_matrix', 'precision_matrix', 'scale_tril']
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    covariance_matrix
    precision_matrix
```

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

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

4.15 LowRankMultivariateNormal

```
class LowRankMultivariateNormal (loc, cov_factor, cov_diag, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution
    arg_constraints = {'cov_diag': <numpyro.distributions.constraints._GreaterThan object>
    support = <numpyro.distributions.constraints._RealVector object>
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
    scale_tril
    covariance_matrix
    precision_matrix
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    entropy ()
```

4.16 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 (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    icdf (q)
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

4.17 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
```

4.18 StudentT

```
class StudentT(df, loc=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(*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

4.19 TruncatedCauchy

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

4.20 TruncatedNormal

class `TruncatedNormal` (*low=0.0, loc=0.0, scale=1.0, validate_args=None*)

Bases: `numpyro.distributions.distribution.TransformedDistribution`

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

reparametrized_params = ['low', 'loc', 'scale']

mean

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

variance

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

4.21 Uniform

class `Uniform` (*low=0.0, high=1.0, validate_args=None*)

Bases: `numpyro.distributions.distribution.TransformedDistribution`

arg_constraints = {'high': `<numpyro.distributions.constraints._Dependent object>`, 'low': `<`

reparametrized_params = ['low', 'high']

mean

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

variance

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

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 (*args, **kwargs)
        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 (*args, **kwargs)
    See numpyro.distributions.distribution.Distribution.log_prob()
mean
    See numpyro.distributions.distribution.Distribution.mean()
variance
    See numpyro.distributions.distribution.Distribution.variance()
```

5.4 BetaBinomial

```
class BetaBinomial (concentration1, concentration0, total_count=1, validate_args=None)
```

Bases: *numpyro.distributions.distribution.Distribution*

Compound distribution comprising of a beta-binomial pair. The probability of success (probs for the Binomial distribution) is unknown and randomly drawn from a Beta distribution prior to a certain number of Bernoulli trials given by total_count.

Parameters

- **concentration1** (*numpy.ndarray*) – 1st concentration parameter (alpha) for the Beta distribution.
- **concentration0** (*numpy.ndarray*) – 2nd concentration parameter (beta) for the Beta distribution.
- **total_count** (*numpy.ndarray*) – number of Bernoulli trials.

```
arg_constraints = {'concentration0': <numpyro.distributions.constraints._GreaterThan
sample (key, sample_shape=())
log_prob (*args, **kwargs)
mean
variance
support
```

5.5 Binomial

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

5.6 BinomialLogits

```
class BinomialLogits (logits, total_count=1, validate_args=None)
```

Bases: *numpyro.distributions.distribution.Distribution*

```
arg_constraints = {'logits': <numpyro.distributions.constraints._Real object>, 'total
```

```

sample (key, sample_shape=())
    See numpyro.distributions.distribution.Distribution.sample()

log_prob (*args, **kwargs)
    See numpyro.distributions.distribution.Distribution.log_prob()

probs

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

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

support

```

5.7 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._Integer object>}

    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()

    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()

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

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

    support

```

5.8 Categorical

```

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

```

5.9 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 (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()

    probs

```

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

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

support

5.10 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 (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
    support
```

5.11 Delta

```
class Delta (value=0.0, log_density=0.0, event_ndim=0, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution

    arg_constraints = {'log_density': <numpyro.distributions.constraints._Real object>, 'value': <numpyro.distributions.constraints._Real object>}
    support = <numpyro.distributions.constraints._Real object>
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
```

5.12 GammaPoisson

```
class GammaPoisson (concentration, rate=1.0, validate_args=None)
    Bases: numpyro.distributions.distribution.Distribution

    Compound distribution comprising of a gamma-poisson pair, also referred to as a gamma-poisson mixture. The
    rate parameter for the Poisson distribution is unknown and randomly drawn from a Gamma distribution.
```

Parameters

- **concentration** (*numpy.ndarray*) – shape parameter (alpha) of the Gamma distribution.
- **rate** (*numpy.ndarray*) – rate parameter (beta) for the Gamma distribution.

```
arg_constraints = {'concentration': <numpyro.distributions.constraints._GreaterThan object>
support = <numpyro.distributions.constraints._IntegerGreaterThan object>
sample (key, sample_shape=())
log_prob (*args, **kwargs)
mean
variance
```

5.13 Multinomial

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

5.14 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._IntegerGreaterThan object>
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
    probs
    mean
        See numpyro.distributions.distribution.Distribution.mean()
    variance
        See numpyro.distributions.distribution.Distribution.variance()
    support
```

5.15 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._IntegerGreaterThan object>
    sample (key, sample_shape=())
        See numpyro.distributions.distribution.Distribution.sample()
    log_prob (*args, **kwargs)
        See numpyro.distributions.distribution.Distribution.log_prob()
```

mean

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

variance

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

support

5.16 OrderedLogistic

class OrderedLogistic (*predictor, cutpoints, validate_args=None*)

Bases: `numpyro.distributions.discrete.CategoricalProbs`

A categorical distribution with ordered outcomes.

References:

1. *Stan Functions Reference*, v2.20 section 12.6, Stan Development Team

Parameters

- **predictor** (`numpy.ndarray`) – prediction in real domain; typically this is output of a linear model.
- **cutpoints** (`numpy.ndarray`) – positions in real domain to separate categories.

arg_constraints = {'cutpoints': `<numpyro.distributions.constraints._OrderedVector object>`}

5.17 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 (**args, **kwargs*)

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

mean

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

variance

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

5.18 PRNGIdentity

class PRNGIdentity

Bases: `numpyro.distributions.distribution.Distribution`

Distribution over `PRNGKey()`. This can be used to draw a batch of `PRNGKey()` using the `seed` handler. Only `sample` method is supported.

```
sample (key, sample_shape=())
```

5.19 ZeroInflatedPoisson

```
class ZeroInflatedPoisson (gate, rate=1.0, validate_args=None)
```

Bases: `numpyro.distributions.distribution.Distribution`

A Zero Inflated Poisson distribution.

Parameters

- **gate** (`numpy.ndarray`) – probability of extra zeros.
- **rate** (`numpy.ndarray`) – rate of Poisson distribution.

```
arg_constraints = {'gate':  <numpyro.distributions.constraints._Interval object>, 'rate':  <numpyro.distributions.constraints._IntegerGreaterThanOrEqualTo object>}
```

```
support = <numpyro.distributions.constraints._IntegerGreaterThanOrEqualTo object>
```

```
sample (key, sample_shape=())
```

```
log_prob (*args, **kwargs)
```

```
mean
```

```
variance
```


6.1 boolean

```
boolean = <numpyro.distributions.constraints._Boolean object>
```

6.2 corr_cholesky

```
corr_cholesky = <numpyro.distributions.constraints._CorrCholesky object>
```

6.3 corr_matrix

```
corr_matrix = <numpyro.distributions.constraints._CorrMatrix object>
```

6.4 dependent

```
dependent = <numpyro.distributions.constraints._Dependent object>
```

6.5 greater_than

```
greater_than(lower_bound)
```

6.6 integer_interval

```
integer_interval(lower_bound, upper_bound)
```

6.7 integer_greater_than

`integer_greater_than(lower_bound)`

6.8 interval

`interval(lower_bound, upper_bound)`

6.9 lower_cholesky

`lower_cholesky = <numpyro.distributions.constraints._LowerCholesky object>`

6.10 multinomial

`multinomial(upper_bound)`

6.11 nonnegative_integer

`nonnegative_integer = <numpyro.distributions.constraints._IntegerGreaterThan object>`

6.12 ordered_vector

`ordered_vector = <numpyro.distributions.constraints._OrderedVector object>`

6.13 positive

`positive = <numpyro.distributions.constraints._GreaterThan object>`

6.14 positive_definite

`positive_definite = <numpyro.distributions.constraints._PositiveDefinite object>`

6.15 positive_integer

`positive_integer = <numpyro.distributions.constraints._IntegerGreaterThan object>`

6.16 real

`real = <numpyro.distributions.constraints._Real object>`

6.17 real_vector

```
real_vector = <numpyro.distributions.constraints._RealVector object>
```

6.18 simplex

```
simplex = <numpyro.distributions.constraints._Simplex object>
```

6.19 unit_interval

```
unit_interval = <numpyro.distributions.constraints._Interval object>
```


7.1 `biject_to`

`biject_to`(*constraint*)

7.2 Transform

```
class Transform
    Bases: object
    domain = <numpyro.distributions.constraints._Real object>
    codomain = <numpyro.distributions.constraints._Real object>
    event_dim = 0
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
    call_with_intermediates(x)
```

7.3 AbsTransform

```
class AbsTransform
    Bases: numpyro.distributions.transforms.Transform
    domain = <numpyro.distributions.constraints._Real object>
    codomain = <numpyro.distributions.constraints._GreaterThan object>
    inv(y)
```

7.4 AffineTransform

```
class AffineTransform(loc, scale, domain=<numpyro.distributions.constraints._Real object>)
    Bases: numpyro.distributions.transforms.Transform

    codomain
    event_dim
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
```

7.5 ComposeTransform

```
class ComposeTransform(parts)
    Bases: numpyro.distributions.transforms.Transform

    domain
    codomain
    event_dim
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
    call_with_intermediates(x)
```

7.6 CorrCholeskyTransform

```
class CorrCholeskyTransform
    Bases: numpyro.distributions.transforms.Transform
```

Transforms a unconstrained real vector x with length $D * (D - 1) / 2$ into the Cholesky factor of a D-dimension correlation matrix. This Cholesky factor is a lower triangular matrix with positive diagonals and unit Euclidean norm for each row. The transform is processed as follows:

1. First we convert x into a lower triangular matrix with the following order:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ x_0 & 1 & 0 & 0 \\ x_1 & x_2 & 1 & 0 \\ x_3 & x_4 & x_5 & 1 \end{bmatrix}$$

2. For each row X_i of the lower triangular part, we apply a *signed* version of class `StickBreakingTransform` to transform X_i into a unit Euclidean length vector using the following steps:
 - a. Scales into the interval $(-1, 1)$ domain: $r_i = \tanh(X_i)$.
 - b. Transforms into an unsigned domain: $z_i = r_i^2$.
 - c. Applies $s_i = \text{StickBreakingTransform}(z_i)$.
 - d. Transforms back into signed domain: $y_i = (\text{sign}(r_i), 1) * \sqrt{s_i}$.

```

domain = <numpyro.distributions.constraints._RealVector object>
codomain = <numpyro.distributions.constraints._CorrCholesky object>
event_dim = 2
inv(y)
log_abs_det_jacobian(x, y, intermediates=None)

```

7.7 ExpTransform

```

class ExpTransform(domain=<numpyro.distributions.constraints._Real object>)
    Bases: numpyro.distributions.transforms.Transform
    codomain
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)

```

7.8 IdentityTransform

```

class IdentityTransform(event_dim=0)
    Bases: numpyro.distributions.transforms.Transform
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)

```

7.9 InvCholeskyTransform

```

class InvCholeskyTransform(domain=<numpyro.distributions.constraints._LowerCholesky object>)
    Bases: numpyro.distributions.transforms.Transform
    Transform via the mapping  $y = x @ x.T$ , where  $x$  is a lower triangular matrix with positive diagonal.
    event_dim = 2
    codomain
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)

```

7.10 LowerCholeskyTransform

```

class LowerCholeskyTransform
    Bases: numpyro.distributions.transforms.Transform
    domain = <numpyro.distributions.constraints._RealVector object>
    codomain = <numpyro.distributions.constraints._LowerCholesky object>
    event_dim = 2

```

```
inv(y)
log_abs_det_jacobian(x, y, intermediates=None)
```

7.11 MultivariateAffineTransform

```
class MultivariateAffineTransform(loc, scale_tril)
    Bases: numpyro.distributions.transforms.Transform
    Transform via the mapping  $y = loc + scale\_tril @ x$ .
```

Parameters

- **loc** – a real vector.
- **scale_tril** – a lower triangular matrix with positive diagonal.

```
domain = <numpyro.distributions.constraints._RealVector object>
codomain = <numpyro.distributions.constraints._RealVector object>
event_dim = 1
inv(y)
log_abs_det_jacobian(x, y, intermediates=None)
```

7.12 OrderedTransform

```
class OrderedTransform
    Bases: numpyro.distributions.transforms.Transform
    Transform a real vector to an ordered vector.
```

References:

1. *Stan Reference Manual v2.20, section 10.6*, Stan Development Team

```
domain = <numpyro.distributions.constraints._RealVector object>
codomain = <numpyro.distributions.constraints._OrderedVector object>
event_dim = 1
inv(y)
log_abs_det_jacobian(x, y, intermediates=None)
```

7.13 PermuteTransform

```
class PermuteTransform(permutation)
    Bases: numpyro.distributions.transforms.Transform
    domain = <numpyro.distributions.constraints._RealVector object>
    codomain = <numpyro.distributions.constraints._RealVector object>
    event_dim = 1
    inv(y)
```

```
log_abs_det_jacobian(x, y, intermediates=None)
```

7.14 PowerTransform

```
class PowerTransform(exponent)
    Bases: numpyro.distributions.transforms.Transform

    domain = <numpyro.distributions.constraints._GreaterThan object>
    codomain = <numpyro.distributions.constraints._GreaterThan object>
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
```

7.15 SigmoidTransform

```
class SigmoidTransform
    Bases: numpyro.distributions.transforms.Transform

    codomain = <numpyro.distributions.constraints._Interval object>
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
```

7.16 StickBreakingTransform

```
class StickBreakingTransform
    Bases: numpyro.distributions.transforms.Transform

    domain = <numpyro.distributions.constraints._RealVector object>
    codomain = <numpyro.distributions.constraints._Simplex object>
    event_dim = 1
    inv(y)
    log_abs_det_jacobian(x, y, intermediates=None)
```


8.1 InverseAutoregressiveTransform

```
class InverseAutoregressiveTransform (autoregressive_nn,          log_scale_min_clip=-5.0,
                                     log_scale_max_clip=3.0)
    Bases: numpyro.distributions.transforms.Transform
```

An implementation of Inverse Autoregressive Flow, using Eq (10) from Kingma et al., 2016,

$$\mathbf{y} = \mu_t + \sigma_t \odot \mathbf{x}$$

where $\mathbf{x}$ are the inputs, $\mathbf{y}$ are the outputs, μ_t, σ_t are calculated from an autoregressive network on $\mathbf{x}$, and $\sigma_t > 0$.

References

1. *Improving Variational Inference with Inverse Autoregressive Flow* [arXiv:1606.04934], Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, Max Welling

```
domain = <numpyro.distributions.constraints._RealVector object>
codomain = <numpyro.distributions.constraints._RealVector object>
event_dim = 1
call_with_intermediates (x)
inv (y)
```

Parameters $\mathbf{y}$ (*numpy.ndarray*) – the output of the transform to be inverted

log_abs_det_jacobian ($x, y, intermediates=None$)
Calculates the elementwise determinant of the log jacobian.

Parameters

- $\mathbf{x}$ (*numpy.ndarray*) – the input to the transform
- $\mathbf{y}$ (*numpy.ndarray*) – the output of the transform

Markov Chain Monte Carlo (MCMC)

9.1 Hamiltonian Monte Carlo

```
class MCMC (sampler,      num_warmup,      num_samples,      num_chains=1,      constrain_fn=None,
             chain_method='parallel', progress_bar=True)
Bases: object
```

Provides access to Markov Chain Monte Carlo inference algorithms in NumPyro.

Note: *chain_method* is an experimental arg, which might be removed in a future version.

Note: Setting *progress_bar=False* will improve the speed for many cases.

Parameters

- **sampler** (*MCMCKernel*) – an instance of `MCMCKernel` that determines the sampler for running MCMC. Currently, only *HMC* and *NUTS* are available.
- **num_warmup** (*int*) – Number of warmup steps.
- **num_samples** (*int*) – Number of samples to generate from the Markov chain.
- **num_chains** (*int*) – Number of Number of MCMC chains to run. By default, chains will be run in parallel using `jax.pmap()`, failing which, chains will be run in sequence.
- **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.
- **chain_method** (*str*) – One of ‘parallel’ (default), ‘sequential’, ‘vectorized’. The method ‘parallel’ is used to execute the drawing process in parallel on XLA devices (CPUs/GPUs/TPUs), If there are not enough devices for ‘parallel’, we fall back to ‘sequential’ method to draw chains sequentially. ‘vectorized’ method is an experimental feature

which vectorizes the drawing method, hence allowing us to collect samples in parallel on a single device.

- **progress_bar** (*bool*) – Whether to enable progress bar updates. Defaults to `True`.

run (*rng_key*, **args*, *extra_fields*=(), *collect_warmup*=*False*, *init_params*=*None*, ***kwargs*)

Run the MCMC samplers and collect samples.

Parameters

- **rng_key** (*random.PRNGKey*) – Random number generator key to be used for the sampling.
- **args** – Arguments to be provided to the `numpyro.infer.mcmc.MCMCKernel.init()` method. These are typically the arguments needed by the *model*.
- **extra_fields** (*tuple or list*) – Extra fields (aside from *z*, *diverging*) from `numpyro.infer.mcmc.HMCState` to collect during the MCMC run.
- **collect_warmup** (*bool*) – Whether to collect samples from the warmup phase. Defaults to *False*.
- **init_params** – Initial parameters to begin sampling. The type must be consistent with the input type to *potential_fn*.
- **kwargs** – Keyword arguments to be provided to the `numpyro.infer.mcmc.MCMCKernel.init()` method. These are typically the keyword arguments needed by the *model*.

get_samples (*group_by_chain*=*False*)

Get samples from the MCMC run.

Parameters **group_by_chain** (*bool*) – Whether to preserve the chain dimension. If *True*, all samples will have `num_chains` as the size of their leading dimension.

Returns Samples having the same data type as *init_params*. The data type is a *dict* keyed on site names if a model containing Pyro primitives is used, but can be any `jaxlib.pytree()`, more generally (e.g. when defining a *potential_fn* for HMC that takes *list* args).

get_extra_fields (*group_by_chain*=*False*)

Get extra fields from the MCMC run.

Parameters **group_by_chain** (*bool*) – Whether to preserve the chain dimension. If *True*, all samples will have `num_chains` as the size of their leading dimension.

Returns Extra fields keyed by field names which are specified in the *extra_fields* keyword of `run()`.

print_summary (*prob*=0.9)

```
class HMC (model=None, potential_fn=None, kinetic_fn=None, step_size=1.0, adapt_step_size=True,
            adapt_mass_matrix=True, dense_mass=False, target_accept_prob=0.8, trajectory_length=6.283185307179586,
            init_strategy=functools.partial(<function
            _init_to_uniform>, radius=2))
```

Bases: `numpyro.infer.mcmc.MCMCKernel`

Hamiltonian Monte Carlo inference, using fixed trajectory length, with provision for step size and mass matrix adaptation.

References:

1. *MCMC Using Hamiltonian Dynamics*, Radford M. Neal

Parameters

- **model** – Python callable containing Pyro *primitives*. If model is provided, *potential_fn* will be inferred using the model.
- **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.
- **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 2π .
- **init_strategy** (*callable*) – a per-site initialization function. See *Initialization Strategies* section for available functions.

init (*rng_key*, *num_warmup*, *init_params=None*, *model_args=()*, *model_kwargs={}*)

sample (*state*)

Run HMC from the given *HMCState* and return the resulting *HMCState*.

Parameters *state* (*HMCState*) – Represents the current state.

Returns Next *state* after running HMC.

class NUTS (*model=None*, *potential_fn=None*, *kinetic_fn=None*, *step_size=1.0*, *adapt_step_size=True*, *adapt_mass_matrix=True*, *dense_mass=False*, *target_accept_prob=0.8*, *trajectory_length=None*, *max_tree_depth=10*, *init_strategy=functools.partial(<function _init_to_uniform>, radius=2)*)

Bases: *numpyro.infer.mcmc.HMC*

Hamiltonian Monte Carlo inference, using the No U-Turn Sampler (NUTS) with adaptive path length and mass matrix adaptation.

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

- **model** – Python callable containing Pyro *primitives*. If model is provided, *potential_fn* will be inferred using the model.

- **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.
- **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. This arg has no effect in NUTS sampler.
- **max_tree_depth** (*int*) – Max depth of the binary tree created during the doubling scheme of NUTS sampler. Defaults to 10.
- **init_strategy** (*callable*) – a per-site initialization function. See [Initialization Strategies](#) section for available functions.

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.

Warning: Instead of using this interface directly, we would highly recommend you to use the higher level `numpyro.infer.MCMC` API instead.

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 = numpyro.sample('beta', dist.Normal(coefs_mean, np.ones(3)))
...     intercept = numpyro.sample('intercept', dist.Normal(0., 10.))
...     return numpyro.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)
>>> init_kernel, sample_kernel = hmc(potential_fn, algo='NUTS')
>>> hmc_state = init_kernel(init_params,
...                         trajectory_length=10,
...                         num_warmup=300)
>>> samples = fori_collect(0, 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, *progbar*=True, *rng_key*=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** (*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 2π .
- **max_tree_depth** (*int*) – Max depth of the binary tree created during the doubling scheme of NUTS sampler. Defaults to 10.
- **progbar** (*bool*) – Whether to enable progress bar updates. Defaults to `True`.
- **rng_key** (*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.infer.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*.
- **energy** - Sum of potential energy and kinetic energy of the current state.
- **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).
- **diverging** - A boolean value to indicate whether the current trajectory is diverging.
- **adapt_state** - A *AdaptState* *namedtuple* which contains adaptation information during warmup:
 - **step_size** - Step size to be used by the integrator in the next iteration.
 - **inverse_mass_matrix** - The inverse mass matrix to be used for the next iteration.
 - **mass_matrix_sqrt** - The square root of mass matrix to be used for the next iteration. In case of dense mass, this is the Cholesky factorization of the mass matrix.
- **rng_key** - random number generator seed used for the iteration.

9.2 MCMC Utilities

initialize_model (*rng_key*, *model*, **model_args*, *init_strategy=functools.partial(<function init_to_uniform>, radius=2)*, ***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_key** (*jax.random.PRNGKey*) – random number generator seed to sample from the prior. The returned *init_params* will have the batch shape *rng_key.shape[: -1]*.
- **model** – Python callable containing Pyro primitives.
- ***model_args** – args provided to the model.
- **init_strategy** (*callable*) – a per-site initialization function. See [Initialization Strategies](#) section for available functions.
- ****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 (*lower*, *upper*, *body_fun*, *init_val*, *transform=<function identity>*, *progressbar=True*, ***progressbar_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, *progressbar=False* will be faster, especially when collecting a lot of samples. Refer to example usage in `hmc()`.

Parameters

- **lower** (*int*) – the index to start the collective work. In other words, we will skip collecting the first *lower* values.
- **upper** (*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 to post-process the values returned by *body_fun*.
- **progressbar** – whether to post progress bar updates.
- ****progressbar_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 *progressbar_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.

consensus (*subposteriors*, *num_draws=None*, *diagonal=False*, *rng_key=None*)

Merges subposteriors following consensus Monte Carlo algorithm.

References:

1. *Bayes and big data: The consensus Monte Carlo algorithm*, Steven L. Scott, Alexander W. Blocker, Fernando V. Bonassi, Hugh A. Chipman, Edward I. George, Robert E. McCulloch

Parameters

- **subposteriors** (*list*) – a list in which each element is a collection of samples.
- **num_draws** (*int*) – number of draws from the merged posterior.
- **diagonal** (*bool*) – whether to compute weights using variance or covariance, defaults to *False* (using covariance).

- **rng_key** (*jax.random.PRNGKey*) – source of the randomness, defaults to *jax.random.PRNGKey(0)*.

Returns if *num_draws* is None, merges subposteriors without resampling; otherwise, returns a collection of *num_draws* samples with the same data structure as each subposterior.

parametric (*subposteriors*, *diagonal=False*)

Merges subposteriors following (embarrassingly parallel) parametric Monte Carlo algorithm.

References:

1. *Asymptotically Exact, Embarrassingly Parallel MCMC*, Willie Neiswanger, Chong Wang, Eric Xing

Parameters

- **subposteriors** (*list*) – a list in which each element is a collection of samples.
- **diagonal** (*bool*) – whether to compute weights using variance or covariance, defaults to *False* (using covariance).

Returns the estimated mean and variance/covariance parameters of the joined posterior

parametric_draws (*subposteriors*, *num_draws*, *diagonal=False*, *rng_key=None*)

Merges subposteriors following (embarrassingly parallel) parametric Monte Carlo algorithm.

References:

1. *Asymptotically Exact, Embarrassingly Parallel MCMC*, Willie Neiswanger, Chong Wang, Eric Xing

Parameters

- **subposteriors** (*list*) – a list in which each element is a collection of samples.
- **num_draws** (*int*) – number of draws from the merged posterior.
- **diagonal** (*bool*) – whether to compute weights using variance or covariance, defaults to *False* (using covariance).
- **rng_key** (*jax.random.PRNGKey*) – source of the randomness, defaults to *jax.random.PRNGKey(0)*.

Returns a collection of *num_draws* samples with the same data structure as each subposterior.

Stochastic Variational Inference (SVI)

class `SVI` (*model*, *guide*, *optim*, *loss*, ****static_kwargs**)

Bases: `object`

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).
- **optim** – an instance of `_NumpyroOptim`.
- **loss** – ELBO loss, i.e. negative Evidence Lower Bound, to minimize.
- **static_kwargs** – static arguments for the model / guide, i.e. arguments that remain constant during fitting.

Returns tuple of (*init_fn*, *update_fn*, *evaluate*).

init (*rng_key*, **args*, ****kwargs**)

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random number generator seed.
- **args** – arguments to the model / guide (these can possibly vary during the course of fitting).
- **kwargs** – keyword arguments to the model / guide (these can possibly vary during the course of fitting).

Returns tuple containing initial `SVIState`, and *get_params*, a callable that transforms unconstrained parameter values from the optimizer to the specified constrained domain

get_params (*svi_state*)

Gets values at *param* sites of the *model* and *guide*.

Parameters *svi_state* – current state of the optimizer.

update (*svi_state*, **args*, ***kwargs*)

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

Parameters

- **svi_state** – current state of SVI.
- **args** – arguments to the model / guide (these can possibly vary during the course of fitting).
- **kwargs** – keyword arguments to the model / guide (these can possibly vary during the course of fitting).

Returns tuple of (*svi_state*, *loss*).

evaluate (*svi_state*, **args*, ***kwargs*)

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

Parameters

- **svi_state** – current state of SVI.
- **args** – arguments to the model / guide (these can possibly vary during the course of fitting).
- **kwargs** – keyword arguments to the model / guide.

Returns evaluate ELBO loss given the current parameter values (held within *svi_state.optim_state*).

10.1 ELBO

class **ELBO** (*num_particles=1*)

Bases: `object`

A trace implementation of ELBO-based SVI. The estimator is constructed along the lines of references [1] and [2]. There are no restrictions on the dependency structure of the model or the guide.

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.

References:

1. *Automated Variational Inference in Probabilistic Programming*, David Wingate, Theo Weber
2. *Black Box Variational Inference*, Rajesh Ranganath, Sean Gerrish, David M. Blei

Parameters **num_particles** – The number of particles/samples used to form the ELBO (gradient) estimators.

loss (*rng_key*, *param_map*, *model*, *guide*, **args*, ***kwargs*)

Evaluates the ELBO with an estimator that uses *num_particles* many samples/particles.

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random number generator seed.
- **param_map** (*dict*) – dictionary of current parameter values keyed by site name.
- **model** – Python callable with NumPyro primitives for the model.

- **guide** – Python callable with NumPyro primitives for the guide.
- **args** – arguments to the model / guide (these can possibly vary during the course of fitting).
- **kwargs** – keyword arguments to the model / guide (these can possibly vary during the course of fitting).

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

10.2 RenyiELBO

class `RenyiELBO` (*alpha=0, num_particles=2*)

Bases: `numpyro.infer.elbo.ELBO`

An implementation of Renyi's α -divergence variational inference following reference [1]. In order for the objective to be a strict lower bound, we require $\alpha \geq 0$. Note, however, that according to reference [1], depending on the dataset $\alpha < 0$ might give better results. In the special case $\alpha = 0$, the objective function is that of the important weighted autoencoder derived in reference [2].

Note: Setting $\alpha < 1$ gives a better bound than the usual ELBO.

Parameters

- **alpha** (*float*) – The order of α -divergence. Here $\alpha \neq 1$. Default is 0.
- **num_particles** – The number of particles/samples used to form the objective (gradient) estimator. Default is 2.

References:

1. *Renyi Divergence Variational Inference*, Yingzhen Li, Richard E. Turner
2. *Importance Weighted Autoencoders*, Yuri Burda, Roger Grosse, Ruslan Salakhutdinov

loss (*rng_key, param_map, model, guide, *args, **kwargs*)

Evaluates the Renyi ELBO with an estimator that uses num_particles many samples/particles.

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random number generator seed.
- **param_map** (*dict*) – dictionary of current parameter values keyed by site name.
- **model** – Python callable with NumPyro primitives for the model.
- **guide** – Python callable with NumPyro primitives for the guide.
- **args** – arguments to the model / guide (these can possibly vary during the course of fitting).
- **kwargs** – keyword arguments to the model / guide (these can possibly vary during the course of fitting).

Returns negative of the Renyi Evidence Lower Bound (ELBO) to be minimized.

Automatic Guide Generation

Warning: The interface for the `contrib.autoguide` module is experimental, and subject to frequent revisions.

11.1 AutoContinuous

```
class AutoContinuous (model,          prefix='auto',          init_strategy=functools.partial(<function
                                _init_to_uniform>, radius=2))
Bases: numpyro.contrib.autoguide.AutoGuide
```

Base class for implementations of continuous-valued Automatic Differentiation Variational Inference [1].

Each derived class implements its own `_get_transform()` method.

Assumes model structure and latent dimension are fixed, and all latent variables are continuous.

Note: We recommend using `AutoContinuousELBO` as the objective function *loss* in *SVI*. In addition, we recommend using `sample_posterior()` method for drawing posterior samples from the autoguide as it will automatically do any unpacking and transformations required to constrain the samples to the support of the latent sites.

Reference:

1. *Automatic Differentiation Variational Inference*, Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, David M. Blei

Parameters

- **model** (*callable*) – A NumPyro model.
- **prefix** (*str*) – a prefix that will be prefixed to all param internal sites.
- **init_strategy** (*callable*) – A per-site initialization function. See *Initialization Strategies* section for available functions.

base_dist

Base distribution of the posterior. By default, it is standard normal.

get_transform (*params*)

Returns the transformation learned by the guide to generate samples from the unconstrained (approximate) posterior.

Parameters *params* (*dict*) – Current parameters of model and autoguide.

Returns the transform of posterior distribution

Return type *Transform*

sample_posterior (*rng_key*, *params*, *sample_shape*=())

Get samples from the learned posterior.

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random key to be used draw samples.
- **params** (*dict*) – Current parameters of model and autoguide.
- **sample_shape** (*tuple*) – batch shape of each latent sample, defaults to ().

Returns a dict containing samples drawn the this guide.

Return type *dict*

median (*params*)

Returns the posterior median value of each latent variable.

Parameters *params* (*dict*) – A dict containing parameter values.

Returns A dict mapping sample site name to median tensor.

Return type *dict*

quantiles (*params*, *quantiles*)

Returns posterior quantiles each latent variable. Example:

```
print(guide.quantiles(opt_state, [0.05, 0.5, 0.95]))
```

Parameters

- **params** (*dict*) – A dict containing parameter values.
- **quantiles** (*list*) – A list of requested quantiles between 0 and 1.

Returns A dict mapping sample site name to a list of quantile values.

Return type *dict*

11.2 AutoDiagonalNormal

```
class AutoDiagonalNormal (model, prefix='auto', init_strategy=functools.partial(<function  
_init_to_uniform>, radius=2))
```

Bases: *numpyro.contrib.autoguide.AutoContinuous*

This implementation of *AutoContinuous* uses a Normal distribution with a diagonal covariance matrix to construct a guide over the entire latent space. The guide does not depend on the model's **args*, ***kwargs*.

Usage:

```
guide = AutoDiagonalNormal(model, ...)
svi = SVI(model, guide, ...)
```

median (*params*)

quantiles (*params*, *quantiles*)

11.3 AutoMultivariateNormal

class AutoMultivariateNormal (*model*, *prefix*='auto', *init_strategy*=*functools.partial*(<function *_init_to_uniform*>, *radius*=2))

Bases: *numpyro.contrib.autoguide.AutoContinuous*

This implementation of *AutoContinuous* uses a MultivariateNormal distribution to construct a guide over the entire latent space. The guide does not depend on the model's **args*, ***kwargs*.

Usage:

```
guide = AutoMultivariateNormal(model, ...)
svi = SVI(model, guide, ...)
```

median (*params*)

quantiles (*params*, *quantiles*)

11.4 AutoIAFNormal

class AutoIAFNormal (*model*, *prefix*='auto', *init_strategy*=*functools.partial*(<function *_init_to_uniform*>, *radius*=2), *num_flows*=3, ***arn_kwargs*)

Bases: *numpyro.contrib.autoguide.AutoContinuous*

This implementation of *AutoContinuous* uses a Diagonal Normal distribution transformed via a InverseAutoregressiveTransform to construct a guide over the entire latent space. The guide does not depend on the model's **args*, ***kwargs*.

Usage:

```
guide = AutoIAFNormal(model, hidden_dims=[20], skip_connections=True, ...)
svi = SVI(model, guide, ...)
```

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random key to be used as the source of randomness to initialize the guide.
- **model** (*callable*) – a generative model.
- **prefix** (*str*) – a prefix that will be prefixed to all param internal sites.
- **init_strategy** (*callable*) – A per-site initialization function.
- **num_flows** (*int*) – the number of flows to be used, defaults to 3.
- ****arn_kwargs** – keywords for constructing autoregressive neural networks, which includes:

- **hidden_dims** (`list[int]`) - the dimensionality of the hidden units per layer. Defaults to `[latent_size, latent_size]`.
- **skip_connections** (`bool`) - whether to add skip connections from the input to the output of each flow. Defaults to `False`.
- **nonlinearity** (`callable`) - the nonlinearity to use in the feedforward network. Defaults to `jax.experimental.stax.Relu()`.

11.5 AutoLaplaceApproximation

class `AutoLaplaceApproximation` (*model*, *prefix*='auto', *init_strategy*=`functools.partial(<function _init_to_uniform>, radius=2)`)
Bases: `numpyro.contrib.autoguide.AutoContinuous`

Laplace approximation (quadratic approximation) approximates the posterior $\log p(z|x)$ by a multivariate normal distribution in the unconstrained space. Under the hood, it uses Delta distributions to construct a MAP guide over the entire (unconstrained) latent space. Its covariance is given by the inverse of the hessian of $-\log p(x, z)$ at the MAP point of z .

Usage:

```
guide = AutoLaplaceApproximation(model, ...)
svi = SVI(model, guide, ...)
```

sample_posterior (*rng_key*, *params*, *sample_shape*=())

get_transform (*params*)

median (*params*)

quantiles (*params*, *quantiles*)

11.6 AutoContinuousELBO

class `AutoContinuousELBO` (*num_particles*=1)
Bases: `numpyro.infer.elbo.ELBO`

An ELBO implementation specific to `AutoContinuous` guides. In those guide, the latent variables of the model are transformed to unconstrained domains. This class provides ELBO of the “transformed” model (i.e. the corresponding model with unconstrained variables) and the guide.

loss (*rng_key*, *param_map*, *model*, *guide*, **args*, ***kwargs*)

Optimizer classes defined here are light wrappers over the corresponding optimizers sourced from `jax.experimental.optimizers` with an interface that is better suited for working with NumPyro inference algorithms.

12.1 Adam

```
class Adam (*args, **kwargs)
    Wrapper class for the JAX optimizer: adam()

get_params (state: Tuple[int, _OptState]) → _Params
    Get current parameter values.

    Parameters state – current optimizer state.

    Returns collection with current value for parameters.

init (params: _Params) → Tuple[int, _OptState]
    Initialize the optimizer with parameters designated to be optimized.

    Parameters params – a collection of numpy arrays.

    Returns initial optimizer state.

update (g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
    Gradient update for the optimizer.

    Parameters
    • g – gradient information for parameters.
    • state – current optimizer state.

    Returns new optimizer state after the update.
```

12.2 Adagrad

```
class Adagrad(*args, **kwargs)
    Wrapper class for the JAX optimizer: adagrad()

    get_params (state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init (params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update (g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
        Gradient update for the optimizer.

        Parameters

        • g – gradient information for parameters.

        • state – current optimizer state.

        Returns new optimizer state after the update.
```

12.3 ClippedAdam

```
class ClippedAdam(*args, clip_norm=10.0, **kwargs)
    Adam optimizer with gradient clipping.

    Parameters clip_norm (float) – All gradient values will be clipped between [-clip_norm, clip_norm].

Reference:

A Method for Stochastic Optimization, Diederik P. Kingma, Jimmy Ba https://arxiv.org/abs/1412.6980

    get_params (state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init (params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update (g, state)
```

12.4 Momentum

```
class Momentum(*args, **kwargs)
    Wrapper class for the JAX optimizer: momentum()

    get_params (state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init (params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update (g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
        Gradient update for the optimizer.

        Parameters
        • g – gradient information for parameters.
        • state – current optimizer state.

        Returns new optimizer state after the update.
```

12.5 RMSProp

```
class RMSProp(*args, **kwargs)
    Wrapper class for the JAX optimizer: rmsprop()

    get_params (state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init (params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update (g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
        Gradient update for the optimizer.

        Parameters
        • g – gradient information for parameters.
        • state – current optimizer state.

        Returns new optimizer state after the update.
```

12.6 RMSPropMomentum

```
class RMSPropMomentum(*args, **kwargs)
    Wrapper class for the JAX optimizer: rmsprop_momentum()

    get_params(state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init(params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update(g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
        Gradient update for the optimizer.

        Parameters
        • g – gradient information for parameters.
        • state – current optimizer state.

        Returns new optimizer state after the update.
```

12.7 SGD

```
class SGD(*args, **kwargs)
    Wrapper class for the JAX optimizer: sgd()

    get_params(state: Tuple[int, _OptState]) → _Params
        Get current parameter values.

        Parameters state – current optimizer state.

        Returns collection with current value for parameters.

    init(params: _Params) → Tuple[int, _OptState]
        Initialize the optimizer with parameters designated to be optimized.

        Parameters params – a collection of numpy arrays.

        Returns initial optimizer state.

    update(g: _Params, state: Tuple[int, _OptState]) → Tuple[int, _OptState]
        Gradient update for the optimizer.

        Parameters
        • g – gradient information for parameters.
        • state – current optimizer state.

        Returns new optimizer state after the update.
```

12.8 SM3

class `SM3` (*args, **kwargs)

Wrapper class for the JAX optimizer: `sm3()`

get_params (state: *Tuple[int, _OptState]*) → *_Params*

Get current parameter values.

Parameters `state` – current optimizer state.

Returns collection with current value for parameters.

init (params: *_Params*) → *Tuple[int, _OptState]*

Initialize the optimizer with parameters designated to be optimized.

Parameters `params` – a collection of numpy arrays.

Returns initial optimizer state.

update (g: *_Params*, state: *Tuple[int, _OptState]*) → *Tuple[int, _OptState]*

Gradient update for the optimizer.

Parameters

- `g` – gradient information for parameters.
- `state` – current optimizer state.

Returns new optimizer state after the update.

This provides a small set of utilities in NumPyro that are used to diagnose posterior samples.

13.1 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*

13.2 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*

13.3 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*

13.4 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 *x.shape*[0] >= 2 and *x.shape*[1] >= 2.

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

Returns R-hat of *x*.

Return type *numpy.ndarray*

13.5 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 *x.shape*[1] >= 4.

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

Returns split R-hat of *x*.

Return type *numpy.ndarray*

13.6 HPDI

hpdi(*x*, *prob*=0.9, *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 x at $(1 - \text{prob}) / 2$ and $(1 + \text{prob}) / 2$.

Return type `numpy.ndarray`

13.7 Summary

print_summary (*samples*, *prob=0.9*, *group_by_chain=True*)

Prints a summary table displaying diagnostics of *samples* from the posterior. The diagnostics displayed are mean, standard deviation, median, the 90% Credibility Interval `hpdi()`, `effective_sample_size()`, and `split_gelman_rubin()`.

Parameters

- **samples** (*dict* or *numpy.ndarray*) – a collection of input samples with left most dimension is chain dimension and second to left most dimension is draw dimension.
- **prob** (*float*) – the probability mass of samples within the HPDI interval.
- **group_by_chain** (*bool*) – If True, each variable in *samples* will be treated as having shape *num_chains* $\times$ *num_samples* $\times$ *sample_shape*. Otherwise, the corresponding shape will be *num_samples* $\times$ *sample_shape* (i.e. without chain dimension).

14.1 `enable_validation`

`enable_validation` (*is_validate=True*)

Enable or disable validation checks in NumPyro. Validation checks provide useful warnings and errors, e.g. NaN checks, validating distribution arguments and support values, etc. which is useful for debugging.

Note: This utility does not take effect under JAX's JIT compilation or vectorized transformation `jax.vmap()`.

Parameters `is_validate` (*bool*) – whether to enable validation checks.

14.2 `validation_enabled`

`validation_enabled` (*is_validate=True*)

Context manager that is useful when temporarily enabling/disabling validation checks.

Parameters `is_validate` (*bool*) – whether to enable validation checks.

14.3 `set_platform`

`set_platform` (*platform=None*)

Changes platform to CPU, GPU, or TPU. This utility only takes effect at the beginning of your program.

Parameters `platform` (*str*) – either 'cpu', 'gpu', or 'tpu'.

14.4 set_host_device_count

`set_host_device_count` (*n*)

By default, XLA considers all CPU cores as one device. This utility tells XLA that there are *n* host (CPU) devices available to use. As a consequence, this allows parallel mapping in JAX `jax.pmap()` to work in CPU platform.

Note: This utility only takes effect at the beginning of your program. Under the hood, this sets the environment variable `XLA_FLAGS=-xla_force_host_platform_device_count=[num_devices]`, where `[num_device]` is the desired number of CPU devices *n*.

Warning: Our understanding of the side effects of using the `xla_force_host_platform_device_count` flag in XLA is incomplete. If you observe some strange phenomenon when using this utility, please let us know through our [issue](#) or forum page. More information is available in this [JAX issue](#).

Parameters *n* (*int*) – number of CPU devices to use.

15.1 Predictive

class Predictive (*model*, *posterior_samples=None*, *guide=None*, *params=None*, *num_samples=None*, *return_sites=None*, *parallel=False*)

Bases: `object`

This class is used to construct predictive distribution. The predictive distribution is obtained by running model conditioned on latent samples from *posterior_samples*.

Warning: The interface for the *Predictive* class is experimental, and might change in the future.

Parameters

- **model** – Python callable containing Pyro primitives.
- **posterior_samples** (*dict*) – dictionary of samples from the posterior.
- **guide** (*callable*) – optional guide to get posterior samples of sites not present in *posterior_samples*.
- **params** (*dict*) – dictionary of values for param sites of model/guide.
- **num_samples** (*int*) – number of samples
- **return_sites** (*list*) – sites to return; by default only sample sites not present in *posterior_samples* are returned.
- **parallel** (*bool*) – whether to predict in parallel using JAX vectorized map `jax.vmap()`. Defaults to False.

Returns dict of samples from the predictive distribution.

get_samples (*rng_key*, **args*, ***kwargs*)

Returns dict of samples from the predictive distribution. By default, only sample sites not contained in

posterior_samples are returned. This can be modified by changing the *return_sites* keyword argument of this *Predictive* instance.

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random key to draw samples.
- **args** – model arguments.
- **kwargs** – model kwargs.

15.2 log_density

log_density (*model, model_args, model_kwargs, params, skip_dist_transforms=False*)

Computes log of joint density for the model given latent values *params*.

Parameters

- **model** – Python callable containing NumPyro primitives.
- **model_args** (*tuple*) – args provided to the model.
- **model_kwargs** (*dict*) – kwargs provided to the model.
- **params** (*dict*) – dictionary of current parameter values keyed by site name.
- **skip_dist_transforms** (*bool*) – whether to compute log probability of a site (if its prior is a transformed distribution) in its base distribution domain.

Returns log of joint density and a corresponding model trace

15.3 transform_fn

transform_fn (*transforms, params, invert=False*)

Callable that applies a transformation from the *transforms* dict to values in the *params* dict and returns the transformed values keyed on the same names.

Parameters

- **transforms** – Dictionary of transforms keyed by names. Names in *transforms* and *params* should align.
- **params** – Dictionary of arrays keyed by names.
- **invert** – Whether to apply the inverse of the transforms.

Returns *dict* of transformed params.

15.4 constrain_fn

constrain_fn (*model, model_args, model_kwargs, transforms, params*)

Gets value at each latent site in *model* given unconstrained parameters *params*. The *transforms* is used to transform these unconstrained parameters to base values of the corresponding priors in *model*. If a prior is a transformed distribution, the corresponding base value lies in the support of base distribution. Otherwise, the base value lies in the support of the distribution.

Parameters

- **model** – a callable containing NumPyro primitives.
- **model_args** (*tuple*) – args provided to the model.
- **model_kwargs** (*dict*) – kwargs provided to the model.
- **transforms** (*dict*) – dictionary of transforms keyed by names. Names in *transforms* and *params* should align.
- **params** (*dict*) – dictionary of unconstrained values keyed by site names.

Returns *dict* of transformed params.

15.5 potential_energy

potential_energy (*model*, *model_args*, *model_kwargs*, *inv_transforms*, *params*)

Computes potential energy of a model given unconstrained params. The *inv_transforms* is used to transform these unconstrained parameters to base values of the corresponding priors in *model*. If a prior is a transformed distribution, the corresponding base value lies in the support of base distribution. Otherwise, the base value lies in the support of the distribution.

Parameters

- **model** – a callable containing NumPyro primitives.
- **model_args** (*tuple*) – args provided to the model.
- **model_kwargs** (*dict*) – kwargs provided to the model.
- **inv_transforms** (*dict*) – dictionary of transforms keyed by names.
- **params** (*dict*) – unconstrained parameters of *model*.

Returns potential energy given unconstrained parameters.

15.6 log_likelihood

log_likelihood (*model*, *posterior_samples*, **args*, ***kwargs*)

Returns log likelihood at observation nodes of model, given samples of all latent variables.

Warning: The interface for the *log_likelihood* function is experimental, and might change in the future.

Parameters

- **model** – Python callable containing Pyro primitives.
- **posterior_samples** (*dict*) – dictionary of samples from the posterior.
- **args** – model arguments.
- **kwargs** – model kwargs.

Returns dict of log likelihoods at observation sites.

15.7 find_valid_initial_params

find_valid_initial_params (*rng_key*, *model*, **model_args*, *init_strategy=functools.partial(<function _init_to_uniform>, radius=2)*, *param_as_improper=False*, *proto-type_params=None*, ***model_kwargs*)

Given a model with Pyro primitives, returns an initial valid unconstrained parameters. This function also returns an *is_valid* flag to say whether the initial parameters are valid.

Parameters

- **rng_key** (*jax.random.PRNGKey*) – random number generator seed to sample from the prior. The returned *init_params* will have the batch shape *rng_key.shape[: -1]*.
- **model** – Python callable containing Pyro primitives.
- ***model_args** – args provided to the model.
- **init_strategy** (*callable*) – a per-site initialization function.
- **param_as_improper** (*bool*) – a flag to decide whether to consider sites with *param* statement as sites with improper priors.
- ****model_kwargs** – kwargs provided to the model.

Returns tuple of (*init_params*, *is_valid*).

15.8 Initialization Strategies

15.8.1 init_to_median

init_to_median (*num_samples=15*)

Initialize to the prior median.

Parameters **num_samples** (*int*) – number of prior points to calculate median.

15.8.2 init_to_prior

init_to_prior ()

Initialize to a prior sample.

15.8.3 init_to_uniform

init_to_uniform (*radius=2*)

Initialize to a random point in the area (*-radius*, *radius*) of unconstrained domain.

Parameters **radius** (*float*) – specifies the range to draw an initial point in the unconstrained domain.

15.8.4 init_to_feasible

init_to_feasible ()

Initialize to an arbitrary feasible point, ignoring distribution parameters.

15.8.5 `init_to_value`

`init_to_value` (*values*)

Initialize to the value specified in *values*. We defer to `init_to_uniform()` strategy for sites which do not appear in *values*.

Parameters **values** (*dict*) – dictionary of initial values keyed by site name.

CHAPTER 16

Indices and tables

- `genindex`
- `search`

n

`numpyro.contrib.autoguide`, 55
`numpyro.diagnostics`, 65
`numpyro.handlers`, 3
`numpyro.infer.util`, 71
`numpyro.optim`, 59
`numpyro.primitives`, 1
`numpyro.util`, 69

A

AbsTransform (class in `numpyro.distributions.transforms`), 35
 Adagrad (class in `numpyro.optim`), 60
 Adam (class in `numpyro.optim`), 59
 AffineTransform (class in `numpyro.distributions.transforms`), 36
 arg_constraints (*BernoulliLogits attribute*), 23
 arg_constraints (*BernoulliProbs attribute*), 23
 arg_constraints (*Beta attribute*), 13
 arg_constraints (*BetaBinomial attribute*), 24
 arg_constraints (*BinomialLogits attribute*), 24
 arg_constraints (*BinomialProbs attribute*), 25
 arg_constraints (*CategoricalLogits attribute*), 25
 arg_constraints (*CategoricalProbs attribute*), 26
 arg_constraints (*Cauchy attribute*), 13
 arg_constraints (*Chi2 attribute*), 14
 arg_constraints (*Delta attribute*), 26
 arg_constraints (*Dirichlet attribute*), 14
 arg_constraints (*Distribution attribute*), 9
 arg_constraints (*Exponential attribute*), 14
 arg_constraints (*Gamma attribute*), 15
 arg_constraints (*GammaPoisson attribute*), 27
 arg_constraints (*GaussianRandomWalk attribute*), 15
 arg_constraints (*HalfCauchy attribute*), 15
 arg_constraints (*HalfNormal attribute*), 16
 arg_constraints (*Independent attribute*), 11
 arg_constraints (*InverseGamma attribute*), 16
 arg_constraints (*LKJ attribute*), 17
 arg_constraints (*LKJCholesky attribute*), 18
 arg_constraints (*LogNormal attribute*), 18
 arg_constraints (*LowRankMultivariateNormal attribute*), 19
 arg_constraints (*MultinomialLogits attribute*), 27
 arg_constraints (*MultinomialProbs attribute*), 27
 arg_constraints (*MultivariateNormal attribute*), 18
 arg_constraints (*Normal attribute*), 19
 arg_constraints (*OrderedLogistic attribute*), 28

arg_constraints (*Pareto attribute*), 20
 arg_constraints (*Poisson attribute*), 28
 arg_constraints (*StudentT attribute*), 20
 arg_constraints (*TransformedDistribution attribute*), 11
 arg_constraints (*TruncatedCauchy attribute*), 20
 arg_constraints (*TruncatedNormal attribute*), 21
 arg_constraints (*Uniform attribute*), 21
 arg_constraints (*ZeroInflatedPoisson attribute*), 29
 AutoContinuous (class in `numpyro.contrib.autoguide`), 55
 AutoContinuousELBO (class in `numpyro.contrib.autoguide`), 58
 autocorrelation() (in module `numpyro.diagnostics`), 65
 autocovariance() (in module `numpyro.diagnostics`), 65
 AutoDiagonalNormal (class in `numpyro.contrib.autoguide`), 56
 AutoIAFNormal (class in `numpyro.contrib.autoguide`), 57
 AutoLaplaceApproximation (class in `numpyro.contrib.autoguide`), 58
 AutoMultivariateNormal (class in `numpyro.contrib.autoguide`), 57

B

base_dist (*AutoContinuous attribute*), 56
 batch_shape (*Distribution attribute*), 9
 Bernoulli() (in module `numpyro.distributions.discrete`), 23
 BernoulliLogits (class in `numpyro.distributions.discrete`), 23
 BernoulliProbs (class in `numpyro.distributions.discrete`), 23
 Beta (class in `numpyro.distributions.continuous`), 13
 BetaBinomial (class in `numpyro.distributions.conjugate`), 24
 biject_to() (in module `numpyro.distributions.transforms`), 35

Binomial() (in module *numpyro.distributions.discrete*), 24
 BinomialLogits (class in *numpyro.distributions.discrete*), 24
 BinomialProbs (class in *numpyro.distributions.discrete*), 25
 block (class in *numpyro.handlers*), 4
 boolean (in module *numpyro.distributions.constraints*), 31

C

call_with_intermediates() (*ComposeTransform* method), 36
 call_with_intermediates() (*InverseAutoregressiveTransform* method), 41
 call_with_intermediates() (*Transform* method), 35
 Categorical() (in module *numpyro.distributions.discrete*), 25
 CategoricalLogits (class in *numpyro.distributions.discrete*), 25
 CategoricalProbs (class in *numpyro.distributions.discrete*), 26
 Cauchy (class in *numpyro.distributions.continuous*), 13
 Chi2 (class in *numpyro.distributions.continuous*), 14
 ClippedAdam (class in *numpyro.optim*), 60
 codomain (*AbsTransform* attribute), 35
 codomain (*AffineTransform* attribute), 36
 codomain (*ComposeTransform* attribute), 36
 codomain (*CorrCholeskyTransform* attribute), 37
 codomain (*ExpTransform* attribute), 37
 codomain (*InvCholeskyTransform* attribute), 37
 codomain (*InverseAutoregressiveTransform* attribute), 41
 codomain (*LowerCholeskyTransform* attribute), 37
 codomain (*MultivariateAffineTransform* attribute), 38
 codomain (*OrderedTransform* attribute), 38
 codomain (*PermuteTransform* attribute), 38
 codomain (*PowerTransform* attribute), 39
 codomain (*SigmoidTransform* attribute), 39
 codomain (*StickBreakingTransform* attribute), 39
 codomain (*Transform* attribute), 35
 ComposeTransform (class in *numpyro.distributions.transforms*), 36
 condition (class in *numpyro.handlers*), 5
 consensus() (in module *numpyro.infer.hmc_util*), 49
 constrain_fn() (in module *numpyro.infer.util*), 72
 corr_cholesky (in module *numpyro.distributions.constraints*), 31
 corr_matrix (in module *numpyro.distributions.constraints*), 31
 CorrCholeskyTransform (class in *numpyro.distributions.transforms*), 36

D

covariance_matrix (*LowRankMultivariateNormal* attribute), 19
 covariance_matrix (*MultivariateNormal* attribute), 18

E

effective_sample_size() (in module *numpyro.diagnostics*), 66
 ELBO (class in *numpyro.infer.elbo*), 52
 enable_validation() (in module *numpyro.distributions.distribution*), 69
 entropy() (*LowRankMultivariateNormal* method), 19
 evaluate() (*SVI* method), 52
 event_dim (*AffineTransform* attribute), 36
 event_dim (*ComposeTransform* attribute), 36
 event_dim (*CorrCholeskyTransform* attribute), 37
 event_dim (*InvCholeskyTransform* attribute), 37
 event_dim (*InverseAutoregressiveTransform* attribute), 41
 event_dim (*LowerCholeskyTransform* attribute), 37
 event_dim (*MultivariateAffineTransform* attribute), 38
 event_dim (*OrderedTransform* attribute), 38
 event_dim (*PermuteTransform* attribute), 38
 event_dim (*StickBreakingTransform* attribute), 39
 event_dim (*Transform* attribute), 35
 event_shape (*Distribution* attribute), 10
 Exponential (class in *numpyro.distributions.continuous*), 14
 ExpTransform (class in *numpyro.distributions.transforms*), 37

F

find_valid_initial_params() (in module *numpyro.infer.util*), 74

`fori_collect()` (in module `numpyro.util`), 49

G

`Gamma` (class in `numpyro.distributions.continuous`), 15
`GammaPoisson` (class in `numpyro.distributions.conjugate`), 26
`GaussianRandomWalk` (class in `numpyro.distributions.continuous`), 15
`gelman_rubin()` (in module `numpyro.diagnostics`), 66
`get_extra_fields()` (MCMC method), 44
`get_params()` (Adagrad method), 60
`get_params()` (Adam method), 59
`get_params()` (ClippedAdam method), 60
`get_params()` (Momentum method), 61
`get_params()` (RMSPProp method), 61
`get_params()` (RMSPPropMomentum method), 62
`get_params()` (SGD method), 62
`get_params()` (SM3 method), 63
`get_params()` (SVI method), 51
`get_samples()` (MCMC method), 44
`get_samples()` (Predictive method), 71
`get_trace()` (trace method), 7
`get_transform()` (AutoContinuous method), 56
`get_transform()` (AutoLaplaceApproximation method), 58
`greater_than()` (in module `numpyro.distributions.constraints`), 31

H

`HalfCauchy` (class in `numpyro.distributions.continuous`), 15
`HalfNormal` (class in `numpyro.distributions.continuous`), 16
`HMC` (class in `numpyro.infer.mcmc`), 44
`hmc()` (in module `numpyro.infer.mcmc`), 46
`HMCState` (in module `numpyro.infer.mcmc`), 48
`hpdi()` (in module `numpyro.diagnostics`), 66

I

`icdf()` (Normal method), 19
`IdentityTransform` (class in `numpyro.distributions.transforms`), 37
`Independent` (class in `numpyro.distributions.distribution`), 11
`init()` (Adagrad method), 60
`init()` (Adam method), 59
`init()` (ClippedAdam method), 60
`init()` (HMC method), 45
`init()` (Momentum method), 61
`init()` (RMSPProp method), 61
`init()` (RMSPPropMomentum method), 62
`init()` (SGD method), 62
`init()` (SM3 method), 63

`init()` (SVI method), 51
`init_kernel()` (in module `numpyro.infer.mcmc.hmc`), 47
`init_to_feasible()` (in module `numpyro.infer.util`), 74
`init_to_median()` (in module `numpyro.infer.util`), 74
`init_to_prior()` (in module `numpyro.infer.util`), 74
`init_to_uniform()` (in module `numpyro.infer.util`), 74
`init_to_value()` (in module `numpyro.infer.util`), 75
`initialize_model()` (in module `numpyro.infer.util`), 48
`integer_greater_than()` (in module `numpyro.distributions.constraints`), 32
`integer_interval()` (in module `numpyro.distributions.constraints`), 31
`interval()` (in module `numpyro.distributions.constraints`), 32
`inv()` (AbsTransform method), 35
`inv()` (AffineTransform method), 36
`inv()` (ComposeTransform method), 36
`inv()` (CorrCholeskyTransform method), 37
`inv()` (ExpTransform method), 37
`inv()` (IdentityTransform method), 37
`inv()` (InvCholeskyTransform method), 37
`inv()` (InverseAutoregressiveTransform method), 41
`inv()` (LowerCholeskyTransform method), 37
`inv()` (MultivariateAffineTransform method), 38
`inv()` (OrderedTransform method), 38
`inv()` (PermuteTransform method), 38
`inv()` (PowerTransform method), 39
`inv()` (SigmoidTransform method), 39
`inv()` (StickBreakingTransform method), 39
`inv()` (Transform method), 35
`InvCholeskyTransform` (class in `numpyro.distributions.transforms`), 37
`InverseAutoregressiveTransform` (class in `numpyro.distributions.flows`), 41
`InverseGamma` (class in `numpyro.distributions.continuous`), 16

L

`LKJ` (class in `numpyro.distributions.continuous`), 17
`LKJCholesky` (class in `numpyro.distributions.continuous`), 17
`log_abs_det_jacobian()` (AffineTransform method), 36
`log_abs_det_jacobian()` (ComposeTransform method), 36
`log_abs_det_jacobian()` (CorrCholeskyTransform method), 37
`log_abs_det_jacobian()` (ExpTransform method), 37

`log_abs_det_jacobian()` (*IdentityTransform method*), 37
`log_abs_det_jacobian()` (*InvCholeskyTransform method*), 37
`log_abs_det_jacobian()` (*InverseAutoregressive-Transform method*), 41
`log_abs_det_jacobian()` (*LowerCholeskyTransform method*), 38
`log_abs_det_jacobian()` (*MultivariateAffine-Transform method*), 38
`log_abs_det_jacobian()` (*OrderedTransform method*), 38
`log_abs_det_jacobian()` (*PermuteTransform method*), 38
`log_abs_det_jacobian()` (*PowerTransform method*), 39
`log_abs_det_jacobian()` (*SigmoidTransform method*), 39
`log_abs_det_jacobian()` (*StickBreakingTransform method*), 39
`log_abs_det_jacobian()` (*Transform method*), 35
`log_density()` (*in module numpyro.infer.util*), 72
`log_likelihood()` (*in module numpyro.infer.util*), 73
`log_prob()` (*BernoulliLogits method*), 23
`log_prob()` (*BernoulliProbs method*), 24
`log_prob()` (*Beta method*), 13
`log_prob()` (*BetaBinomial method*), 24
`log_prob()` (*BinomialLogits method*), 25
`log_prob()` (*BinomialProbs method*), 25
`log_prob()` (*CategoricalLogits method*), 25
`log_prob()` (*CategoricalProbs method*), 26
`log_prob()` (*Cauchy method*), 13
`log_prob()` (*Delta method*), 26
`log_prob()` (*Dirichlet method*), 14
`log_prob()` (*Distribution method*), 10
`log_prob()` (*Exponential method*), 14
`log_prob()` (*Gamma method*), 15
`log_prob()` (*GammaPoisson method*), 27
`log_prob()` (*GaussianRandomWalk method*), 15
`log_prob()` (*HalfCauchy method*), 16
`log_prob()` (*HalfNormal method*), 16
`log_prob()` (*Independent method*), 11
`log_prob()` (*LKJCholesky method*), 18
`log_prob()` (*LowRankMultivariateNormal method*), 19
`log_prob()` (*MultinomialLogits method*), 27
`log_prob()` (*MultinomialProbs method*), 27
`log_prob()` (*MultivariateNormal method*), 18
`log_prob()` (*Normal method*), 19
`log_prob()` (*Poisson method*), 28
`log_prob()` (*StudentT method*), 20
`log_prob()` (*TransformedDistribution method*), 12
`log_prob()` (*ZeroInflatedPoisson method*), 29
`LogNormal` (*class in numpyro.distributions.continuous*), 18
`loss()` (*AutoContinuousELBO method*), 58
`loss()` (*ELBO method*), 52
`loss()` (*RenyiELBO method*), 53
`lower_cholesky` (*in module numpyro.distributions.constraints*), 32
`LowerCholeskyTransform` (*class in numpyro.distributions.transforms*), 37
`LowRankMultivariateNormal` (*class in numpyro.distributions.continuous*), 19

M

`MCMC` (*class in numpyro.infer.mcmc*), 43
`mean` (*BernoulliLogits attribute*), 23
`mean` (*BernoulliProbs attribute*), 24
`mean` (*Beta attribute*), 13
`mean` (*BetaBinomial attribute*), 24
`mean` (*BinomialLogits attribute*), 25
`mean` (*BinomialProbs attribute*), 25
`mean` (*CategoricalLogits attribute*), 25
`mean` (*CategoricalProbs attribute*), 26
`mean` (*Cauchy attribute*), 14
`mean` (*Delta attribute*), 26
`mean` (*Dirichlet attribute*), 14
`mean` (*Distribution attribute*), 10
`mean` (*Exponential attribute*), 14
`mean` (*Gamma attribute*), 15
`mean` (*GammaPoisson attribute*), 27
`mean` (*GaussianRandomWalk attribute*), 15
`mean` (*HalfCauchy attribute*), 16
`mean` (*HalfNormal attribute*), 16
`mean` (*Independent attribute*), 11
`mean` (*InverseGamma attribute*), 16
`mean` (*LKJ attribute*), 17
`mean` (*LogNormal attribute*), 18
`mean` (*LowRankMultivariateNormal attribute*), 19
`mean` (*MultinomialLogits attribute*), 27
`mean` (*MultinomialProbs attribute*), 27
`mean` (*MultivariateNormal attribute*), 18
`mean` (*Normal attribute*), 19
`mean` (*Pareto attribute*), 20
`mean` (*Poisson attribute*), 28
`mean` (*StudentT attribute*), 20
`mean` (*TransformedDistribution attribute*), 12
`mean` (*TruncatedCauchy attribute*), 20
`mean` (*TruncatedNormal attribute*), 21
`mean` (*Uniform attribute*), 21
`mean` (*ZeroInflatedPoisson attribute*), 29
`median()` (*AutoContinuous method*), 56
`median()` (*AutoDiagonalNormal method*), 57
`median()` (*AutoLaplaceApproximation method*), 58
`median()` (*AutoMultivariateNormal method*), 57
`module()` (*in module numpyro.primitives*), 2

Momentum (class in `numpyro.optim`), 61
`multinomial()` (in module `numpyro.distributions.constraints`), 32
`Multinomial()` (in module `numpyro.distributions.discrete`), 27
`MultinomialLogits` (class in `numpyro.distributions.discrete`), 27
`MultinomialProbs` (class in `numpyro.distributions.discrete`), 27
`MultivariateAffineTransform` (class in `numpyro.distributions.transforms`), 38
`MultivariateNormal` (class in `numpyro.distributions.continuous`), 18

N

`nonnegative_integer` (in module `numpyro.distributions.constraints`), 32
`Normal` (class in `numpyro.distributions.continuous`), 19
`numpyro.contrib.autoguide` (module), 55
`numpyro.diagnostics` (module), 65
`numpyro.handlers` (module), 3
`numpyro.infer.util` (module), 71
`numpyro.optim` (module), 59
`numpyro.primitives` (module), 1
`numpyro.util` (module), 69
`NUTS` (class in `numpyro.infer.mcmc`), 45

O

`ordered_vector` (in module `numpyro.distributions.constraints`), 32
`OrderedLogistic` (class in `numpyro.distributions.discrete`), 28
`OrderedTransform` (class in `numpyro.distributions.transforms`), 38

P

`param()` (in module `numpyro.primitives`), 1
`parametric()` (in module `numpyro.infer.hmc_util`), 50
`parametric_draws()` (in module `numpyro.infer.hmc_util`), 50
`Pareto` (class in `numpyro.distributions.continuous`), 20
`PermuteTransform` (class in `numpyro.distributions.transforms`), 38
`plate` (class in `numpyro.primitives`), 2
`Poisson` (class in `numpyro.distributions.discrete`), 28
`positive` (in module `numpyro.distributions.constraints`), 32
`positive_definite` (in module `numpyro.distributions.constraints`), 32
`positive_integer` (in module `numpyro.distributions.constraints`), 32
`postprocess_message()` (trace method), 7
`potential_energy()` (in module `numpyro.infer.util`), 73

`PowerTransform` (class in `numpyro.distributions.transforms`), 39
`precision_matrix` (`LowRankMultivariateNormal` attribute), 19
`precision_matrix` (`MultivariateNormal` attribute), 18
`Predictive` (class in `numpyro.infer.util`), 71
`print_summary()` (in module `numpyro.diagnostics`), 67
`print_summary()` (MCMC method), 44
`PRNGIdentity` (class in `numpyro.distributions.discrete`), 28
`probs` (`BernoulliLogits` attribute), 23
`probs` (`BinomialLogits` attribute), 25
`probs` (`CategoricalLogits` attribute), 25
`probs` (`MultinomialLogits` attribute), 27
`process_message()` (block method), 4
`process_message()` (condition method), 5
`process_message()` (replay method), 5
`process_message()` (scale method), 6
`process_message()` (seed method), 6
`process_message()` (substitute method), 7

Q

`quantiles()` (`AutoContinuous` method), 56
`quantiles()` (`AutoDiagonalNormal` method), 57
`quantiles()` (`AutoLaplaceApproximation` method), 58
`quantiles()` (`AutoMultivariateNormal` method), 57

R

`real` (in module `numpyro.distributions.constraints`), 32
`real_vector` (in module `numpyro.distributions.constraints`), 33
`RenyiELBO` (class in `numpyro.infer.elbo`), 53
`reparameterized_params` (`Independent` attribute), 11
`reparameterized_params` (`Cauchy` attribute), 13
`reparameterized_params` (`Distribution` attribute), 9
`reparameterized_params` (`Exponential` attribute), 14
`reparameterized_params` (`Gamma` attribute), 15
`reparameterized_params` (`GaussianRandomWalk` attribute), 15
`reparameterized_params` (`HalfCauchy` attribute), 15
`reparameterized_params` (`HalfNormal` attribute), 16
`reparameterized_params` (`InverseGamma` attribute), 16
`reparameterized_params` (`LogNormal` attribute), 18
`reparameterized_params` (`MultivariateNormal` attribute), 18

reparametrized_params (*Normal attribute*), 19
 reparametrized_params (*StudentT attribute*), 20
 reparametrized_params (*TruncatedCauchy attribute*), 20
 reparametrized_params (*TruncatedNormal attribute*), 21
 reparametrized_params (*Uniform attribute*), 21
 replay (*class in numpyro.handlers*), 5
 RMSProp (*class in numpyro.optim*), 61
 RMSPropMomentum (*class in numpyro.optim*), 62
 run() (*MCMC method*), 44

S

sample() (*BernoulliLogits method*), 23
 sample() (*BernoulliProbs method*), 24
 sample() (*Beta method*), 13
 sample() (*BetaBinomial method*), 24
 sample() (*BinomialLogits method*), 24
 sample() (*BinomialProbs method*), 25
 sample() (*CategoricalLogits method*), 25
 sample() (*CategoricalProbs method*), 26
 sample() (*Cauchy method*), 13
 sample() (*Delta method*), 26
 sample() (*Dirichlet method*), 14
 sample() (*Distribution method*), 10
 sample() (*Exponential method*), 14
 sample() (*Gamma method*), 15
 sample() (*GammaPoisson method*), 27
 sample() (*GaussianRandomWalk method*), 15
 sample() (*HalfCauchy method*), 16
 sample() (*HalfNormal method*), 16
 sample() (*HMC method*), 45
 sample() (*in module numpyro.primitives*), 1
 sample() (*Independent method*), 11
 sample() (*LKJCholesky method*), 18
 sample() (*LowRankMultivariateNormal method*), 19
 sample() (*MultinomialLogits method*), 27
 sample() (*MultinomialProbs method*), 27
 sample() (*MultivariateNormal method*), 18
 sample() (*Normal method*), 19
 sample() (*Poisson method*), 28
 sample() (*PRNGIdentity method*), 28
 sample() (*StudentT method*), 20
 sample() (*TransformedDistribution method*), 12
 sample() (*ZeroInflatedPoisson method*), 29
 sample_kernel() (*in module numpyro.infer.mcmc.hmc*), 48
 sample_posterior() (*AutoContinuous method*), 56
 sample_posterior() (*AutoLaplaceApproximation method*), 58
 sample_with_intermediates() (*Distribution method*), 10
 sample_with_intermediates() (*Transformed-Distribution method*), 12

scale (*class in numpyro.handlers*), 6
 scale_tril (*LowRankMultivariateNormal attribute*), 19
 seed (*class in numpyro.handlers*), 6
 set_default_validate_args() (*Distribution static method*), 9
 set_host_device_count() (*in module numpyro.util*), 70
 set_platform() (*in module numpyro.util*), 69
 SGD (*class in numpyro.optim*), 62
 SigmoidTransform (*class in numpyro.distributions.transforms*), 39
 simplex (*in module numpyro.distributions.constraints*), 33
 SM3 (*class in numpyro.optim*), 63
 split_gelman_rubin() (*in module numpyro.diagnostics*), 66
 StickBreakingTransform (*class in numpyro.distributions.transforms*), 39
 StudentT (*class in numpyro.distributions.continuous*), 20
 substitute (*class in numpyro.handlers*), 6
 support (*BernoulliLogits attribute*), 23
 support (*BernoulliProbs attribute*), 24
 support (*Beta attribute*), 13
 support (*BetaBinomial attribute*), 24
 support (*BinomialLogits attribute*), 25
 support (*BinomialProbs attribute*), 25
 support (*CategoricalLogits attribute*), 26
 support (*CategoricalProbs attribute*), 26
 support (*Cauchy attribute*), 13
 support (*Delta attribute*), 26
 support (*Dirichlet attribute*), 14
 support (*Distribution attribute*), 9
 support (*Exponential attribute*), 14
 support (*Gamma attribute*), 15
 support (*GammaPoisson attribute*), 27
 support (*GaussianRandomWalk attribute*), 15
 support (*HalfCauchy attribute*), 15
 support (*HalfNormal attribute*), 16
 support (*Independent attribute*), 11
 support (*InverseGamma attribute*), 16
 support (*LKJ attribute*), 17
 support (*LKJCholesky attribute*), 18
 support (*LowRankMultivariateNormal attribute*), 19
 support (*MultinomialLogits attribute*), 27
 support (*MultinomialProbs attribute*), 28
 support (*MultivariateNormal attribute*), 18
 support (*Normal attribute*), 19
 support (*Pareto attribute*), 20
 support (*Poisson attribute*), 28
 support (*StudentT attribute*), 20
 support (*TransformedDistribution attribute*), 12
 support (*ZeroInflatedPoisson attribute*), 29

SVI (*class in `numpyro.infer.svi`*), 51

T

`to_event()` (*Distribution method*), 10

`trace` (*class in `numpyro.handlers`*), 7

`Transform` (*class in `numpyro.distributions.transforms`*), 35

`transform_fn()` (*in module `numpyro.infer.util`*), 72

`transform_with_intermediates()` (*Distribution method*), 10

`transform_with_intermediates()` (*TransformedDistribution method*), 12

`TransformedDistribution` (*class in `numpyro.distributions.distribution`*), 11

`TruncatedCauchy` (*class in `numpyro.distributions.continuous`*), 20

`TruncatedNormal` (*class in `numpyro.distributions.continuous`*), 21

U

`Uniform` (*class in `numpyro.distributions.continuous`*), 21

`unit_interval` (*in module `numpyro.distributions.constraints`*), 33

`update()` (*Adagrad method*), 60

`update()` (*Adam method*), 59

`update()` (*ClippedAdam method*), 60

`update()` (*Momentum method*), 61

`update()` (*RMSPProp method*), 61

`update()` (*RMSPPropMomentum method*), 62

`update()` (*SGD method*), 62

`update()` (*SM3 method*), 63

`update()` (*SVI method*), 51

V

`validation_enabled()` (*in module `numpyro.distributions.distribution`*), 69

`variance` (*BernoulliLogits attribute*), 23

`variance` (*BernoulliProbs attribute*), 24

`variance` (*Beta attribute*), 13

`variance` (*BetaBinomial attribute*), 24

`variance` (*BinomialLogits attribute*), 25

`variance` (*BinomialProbs attribute*), 25

`variance` (*CategoricalLogits attribute*), 26

`variance` (*CategoricalProbs attribute*), 26

`variance` (*Cauchy attribute*), 14

`variance` (*Delta attribute*), 26

`variance` (*Dirichlet attribute*), 14

`variance` (*Distribution attribute*), 10

`variance` (*Exponential attribute*), 14

`variance` (*Gamma attribute*), 15

`variance` (*GammaPoisson attribute*), 27

`variance` (*GaussianRandomWalk attribute*), 15

`variance` (*HalfCauchy attribute*), 16

`variance` (*HalfNormal attribute*), 16

`variance` (*Independent attribute*), 11

`variance` (*InverseGamma attribute*), 16

`variance` (*LKJ attribute*), 17

`variance` (*LogNormal attribute*), 18

`variance` (*LowRankMultivariateNormal attribute*), 19

`variance` (*MultinomialLogits attribute*), 27

`variance` (*MultinomialProbs attribute*), 28

`variance` (*MultivariateNormal attribute*), 19

`variance` (*Normal attribute*), 19

`variance` (*Pareto attribute*), 20

`variance` (*Poisson attribute*), 28

`variance` (*StudentT attribute*), 20

`variance` (*TransformedDistribution attribute*), 12

`variance` (*TruncatedCauchy attribute*), 20

`variance` (*TruncatedNormal attribute*), 21

`variance` (*Uniform attribute*), 21

`variance` (*ZeroInflatedPoisson attribute*), 29

Z

`ZeroInflatedPoisson` (*class in `numpyro.distributions.discrete`*), 29