
Pyro Documentation

Release 0.2.0-1-gf5b7e6d

Uber AI Labs

May 22, 2018

Contents:

1	Installation	1
1.1	Install from Source	1
2	Getting Started	3
3	Primitives	5
4	Inference	9
4.1	SVI	9
4.2	ELBO	10
4.3	Importance	12
4.4	Inference Utilities	12
4.5	MCMC	13
5	Distributions	17
5.1	PyTorch Distributions	17
5.2	Pyro Distributions	20
5.3	Transformed Distributions	28
6	Parameters	31
6.1	ParamStore	31
7	Neural Network	33
7.1	AutoRegressiveNN	33
8	Optimization	35
8.1	Pyro Optimizers	35
8.2	Pytorch Optimizers	36
9	Poutine (Effect handlers)	37
9.1	Handlers	37
9.2	Trace	42
9.3	Messengers	44
9.4	Runtime	48
9.5	Utilities	48
10	Contributed Code	51
10.1	Automatic Guide Generation	51

10.2	Gaussian Processes	54
10.3	Named Data Structures	78
11	Indices and tables	81
	Python Module Index	83

1.1 Install from Source

Pyro supports Python 2.7.* and Python 3. To setup, install [PyTorch](#) then run:

```
pip install pyro-ppl
```

or install from source:

```
git clone https://github.com/uber/pyro.git
cd pyro
python setup.py install
```


CHAPTER 2

Getting Started

- [Install Pyro.](#)
- [Learn the basic concepts of Pyro: models and inference.](#)
- [Dive in to other tutorials and examples.](#)

sample (*name*, *fn*, **args*, ***kwargs*)

Calls the stochastic function *fn* with additional side-effects depending on *name* and the enclosing context (e.g. an inference algorithm). See [Intro I](#) and [Intro II](#) for a discussion.

Parameters

- **name** – name of sample
- **fn** – distribution class or function
- **obs** – observed datum (optional; should only be used in context of inference) optionally specified in *kwargs*
- **infer** (*dict*) – Optional dictionary of inference parameters specified in *kwargs*. See inference documentation for details.

Returns sample

param (*name*, **args*, ***kwargs*)

Saves the variable as a parameter in the param store. To interact with the param store or write to disk, see [Parameters](#).

Parameters **name** – name of parameter

Returns parameter

module (*name*, *nn_module*, *update_module_params=False*)

Takes a `torch.nn.Module` and registers its parameters with the `ParamStore`. In conjunction with the `ParamStore` `save()` and `load()` functionality, this allows the user to save and load modules.

Parameters

- **name** (*str*) – name of module
- **nn_module** (*torch.nn.Module*) – the module to be registered with Pyro
- **update_module_params** – determines whether Parameters in the PyTorch module get overridden with the values found in the `ParamStore` (if any). Defaults to *False*

Returns `torch.nn.Module`

random_module (*name*, *nn_module*, *prior*, **args*, ***kwargs*)

Places a prior over the parameters of the module *nn_module*. Returns a distribution (callable) over *nn.Module*'s, which upon calling returns a sampled *nn.Module*.

See the [Bayesian Regression](#) tutorial for an example.

Parameters

- **name** (*str*) – name of pyro module
- **nn_module** (*torch.nn.Module*) – the module to be registered with pyro
- **prior** – pyro distribution, stochastic function, or python dict with parameter names as keys and respective distributions/stochastic functions as values.

Returns a callable which returns a sampled module

class irange (*name*, *size*, *subsample_size*=None, *subsample*=None, *use_cuda*=None)

Non-vectorized version of *iarange*. See *iarange* for details.

Parameters

- **name** (*str*) – A name that will be used for this site in a Trace.
- **size** (*int*) – The size of the collection being subsampled (like `stop` in builtin `range()`).
- **subsample_size** (*int*) – Size of minibatches used in subsampling. Defaults to *size*.
- **subsample** (Anything supporting `len()`.) – Optional custom subsample for user-defined subsampling schemes. If specified, then *subsample_size* will be set to `len(subsample)`.
- **use_cuda** (*bool*) – Optional bool specifying whether to use cuda tensors for internal `log_prob` computations. Defaults to `torch.Tensor.is_cuda`.

Returns A reusable iterator yielding a sequence of integers.

Examples:

```
>>> for i in irange('data', 100, subsample_size=10):
    if z[i]: # Prevents vectorization.
        observe('obs_{}'.format(i), normal, data[i], loc, scale)
```

See [SVI Part II](#) for an extended discussion.

class iarange (*name*, *size*=None, *subsample_size*=None, *subsample*=None, *dim*=None, *use_cuda*=None)

Context manager for conditionally independent ranges of variables.

iarange is similar to `torch.arange()` in that it yields an array of indices by which other tensors can be indexed. *iarange* differs from `torch.arange()` in that it also informs inference algorithms that the variables being indexed are conditionally independent. To do this, *iarange* is provided as context manager rather than a function, and users must guarantee that all computation within an *iarange* context is conditionally independent:

```
with iarange("name", size) as ind:
    # ...do conditionally independent stuff with ind...
```

Additionally, *iarange* can take advantage of the conditional independence assumptions by subsampling the indices and informing inference algorithms to scale various computed values. This is typically used to subsample minibatches of data:

```
with iarange("data", len(data), subsample_size=100) as ind:
    batch = data[ind]
    assert len(batch) == 100
```

By default `subsample_size=False` and this simply yields a `torch.arange(0, size)`. If $0 < \text{subsample_size} \leq \text{size}$ this yields a single random batch of indices of size `subsample_size` and scales all log likelihood terms by `size/batch_size`, within this context.

Warning: This is only correct if all computation is conditionally independent within the context.

Parameters

- **name** (*str*) – A unique name to help inference algorithms match `iarange` sites between models and guides.
- **size** (*int*) – Optional size of the collection being subsampled (like `stop` in builtin `range`).
- **subsample_size** (*int*) – Size of minibatches used in subsampling. Defaults to `size`.
- **subsample** (Anything supporting `len()`) – Optional custom subsample for user-defined subsampling schemes. If specified, then `subsample_size` will be set to `len(subsample)`.
- **dim** (*int*) – An optional dimension to use for this independence index. If specified, `dim` should be negative, i.e. should index from the right. If not specified, `dim` is set to the rightmost dim that is left of all enclosing `iarange` contexts.
- **use_cuda** (*bool*) – Optional bool specifying whether to use cuda tensors for `subsample` and `log_prob`. Defaults to `torch.Tensor.is_cuda`.

Returns A reusable context manager yielding a single 1-dimensional `torch.Tensor` of indices.

Examples:

```
# This version simply declares independence:
>>> with iarange('data'):
    sample('obs', Normal(loc, scale), obs=data)

# This version subsamples data in vectorized way:
>>> with iarange('data', 100, subsample_size=10) as ind:
    sample('obs', Normal(loc, scale), obs=data[ind])

# This wraps a user-defined subsampling method for use in pyro:
>>> ind = my_custom_subsample
>>> with iarange('data', 100, subsample=ind):
    sample('obs', Normal(loc, scale), obs=data[ind])

# This reuses two different independence contexts.
>>> x_axis = iarange('outer', 320, dim=-1)
>>> y_axis = iarange('inner', 200, dim=-2)
>>> with x_axis:
    x_noise = sample("x_noise", Normal(loc, scale).expand_by([320]))
>>> with y_axis:
    y_noise = sample("y_noise", Normal(loc, scale).expand_by([200, 1]))
>>> with x_axis, y_axis:
    xy_noise = sample("xy_noise", Normal(loc, scale).expand_by([200, 320]))
```

See [SVI Part II](#) for an extended discussion.

get_param_store()

Returns the ParamStore

clear_param_store()

Clears the ParamStore. This is especially useful if you're working in a REPL.

validation_enabled(*args, **kws)

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

Parameters **is_validate** (*bool*) – (optional; defaults to True) temporary validation check override.

enable_validation(is_validate=True)

Enable or disable validation checks in Pyro. Validation checks provide useful warnings and errors, e.g. NaN checks, validating distribution arguments and support values, etc. which is useful for debugging. Since some of these checks may be expensive, we recommend turning this off for mature models.

Parameters **is_validate** (*bool*) – (optional; defaults to True) whether to enable validation checks.

In the context of probabilistic modeling, learning is usually called inference. In the particular case of Bayesian inference, this often involves computing (approximate) posterior distributions. In the case of parameterized models, this usually involves some sort of optimization. Pyro supports multiple inference algorithms, with support for stochastic variational inference (SVI) being the most extensive. Look here for more inference algorithms in future versions of Pyro.

See [Intro II](#) for a discussion of inference in Pyro.

4.1 SVI

class `SVI` (*model*, *guide*, *optim*, *loss*, *loss_and_grads=None*, ***kwargs*)

Bases: `object`

Parameters

- **model** – the model (callable containing Pyro primitives)
- **guide** – the guide (callable containing Pyro primitives)
- **optim** (`pyro.optim.PyroOptim`) – a wrapper for a PyTorch optimizer
- **loss** (`pyro.infer.elbo.ELBO`) – an instance of a subclass of `ELBO`. Pyro provides three built-in losses: `Trace_ELBO`, `Trace_ELBO`, and `Trace_ELBO`. See the `ELBO` docs to learn how to implement a custom loss.

A unified interface for stochastic variational inference in Pyro. The most commonly used loss is `loss=Trace_ELBO()`. See the tutorial [SVI Part I](#) for a discussion.

evaluate_loss (**args*, ***kwargs*)

Returns estimate of the loss

Return type `float`

Evaluate the loss function. Any args or kwargs are passed to the model and guide.

step (*args, **kwargs)

Returns estimate of the loss

Return type float

Take a gradient step on the loss function (and any auxiliary loss functions generated under the hood by *loss_and_grads*). Any args or kwargs are passed to the model and guide

4.2 ELBO

class ELBO (num_particles=1, max_iarange_nesting=inf, strict_enumeration_warning=True)

Bases: object

ELBO is the top-level interface for stochastic variational inference via optimization of the evidence lower bound. Most users will not interact with *ELBO* directly; instead they will interact with *SVI*. *ELBO* dispatches to *Trace_ELBO* and *TraceGraph_ELBO*, where the internal implementations live.

Parameters

- **num_particles** – The number of particles/samples used to form the ELBO (gradient) estimators.
- **max_iarange_nesting** (*int*) – Optional bound on max number of nested *pyro.iarange()* contexts. This is only required to enumerate over sample sites in parallel, e.g. if a site sets `infer={"enumerate": "parallel"}`.
- **strict_enumeration_warning** (*bool*) – Whether to warn about possible misuse of enumeration, i.e. that `pyro.infer.traceenum_elbo.TraceEnum_ELBO` is used iff there are enumerated sample sites.

References

[1] *Automated Variational Inference in Probabilistic Programming* David Wingate, Theo Weber

[2] *Black Box Variational Inference*, Rajesh Ranganath, Sean Gerrish, David M. Blei

class Trace_ELBO (num_particles=1, max_iarange_nesting=inf, strict_enumeration_warning=True)

Bases: `pyro.infer.elbo.ELBO`

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. The gradient estimator includes partial Rao-Blackwellization for reducing the variance of the estimator when non-reparameterizable random variables are present. The Rao-Blackwellization is partial in that it only uses conditional independence information that is marked by *iarange* contexts. For more fine-grained Rao-Blackwellization, see *TraceGraph_ELBO*.

References

[1] **Automated Variational Inference in Probabilistic Programming**, David Wingate, Theo Weber

[2] **Black Box Variational Inference**, Rajesh Ranganath, Sean Gerrish, David M. Blei

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

Returns returns an estimate of the ELBO

Return type float

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

loss_and_grads (model, guide, *args, **kwargs)

Returns returns an estimate of the ELBO

Return type float

Computes the ELBO as well as the surrogate ELBO that is used to form the gradient estimator. Performs backward on the latter. Num_particle many samples are used to form the estimators.

```
class TraceGraph_ELBO (num_particles=1,                               max_iarange_nesting=inf,
                        strict_enumeration_warning=True)
```

Bases: `pyro.infer.elbo.ELBO`

A TraceGraph implementation of ELBO-based SVI. The gradient estimator is constructed along the lines of reference [1] specialized to the case of the ELBO. It supports arbitrary dependency structure for the model and guide as well as baselines for non-reparameterizable random variables. Where possible, conditional dependency information as recorded in the `Trace` is used to reduce the variance of the gradient estimator. In particular three kinds of conditional dependency information are used to reduce variance: - the sequential order of samples (z is sampled after y => y does not depend on z) - `iarange` generators - `irange` generators

References

[1] *Gradient Estimation Using Stochastic Computation Graphs*, John Schulman, Nicolas Heess, Theophane Weber, Pieter Abbeel

[2] *Neural Variational Inference and Learning in Belief Networks* Andriy Mnih, Karol Gregor

```
loss (model, guide, *args, **kwargs)
```

Returns returns an estimate of the ELBO

Return type float

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

```
loss_and_grads (model, guide, *args, **kwargs)
```

Returns returns an estimate of the ELBO

Return type float

Computes the ELBO as well as the surrogate ELBO that is used to form the gradient estimator. Performs backward on the latter. Num_particle many samples are used to form the estimators. If baselines are present, a baseline loss is also constructed and differentiated.

```
class TraceEnum_ELBO (num_particles=1, max_iarange_nesting=inf, strict_enumeration_warning=True)
```

Bases: `pyro.infer.elbo.ELBO`

A trace implementation of ELBO-based SVI that supports enumeration over discrete sample sites.

To enumerate over a sample site, the guide's sample site must specify either `infer={'enumerate': 'sequential'}` or `infer={'enumerate': 'parallel'}`. To configure all sites at once, use `config_enumerate`()`.

This assumes restricted dependency structure on the model and guide: variables outside of an `iarange` can never depend on variables inside that `iarange`.

```
loss (model, guide, *args, **kwargs)
```

Returns returns an estimate of the ELBO

Return type float

Estimates the ELBO using num_particles many samples (particles).

```
loss_and_grads (model, guide, *args, **kwargs)
```

Returns returns an estimate of the ELBO

Return type `float`

Estimates the ELBO using `num_particles` many samples (particles). Performs backward on the ELBO of each particle.

4.3 Importance

class `Importance` (*model*, *guide=None*, *num_samples=None*)

Bases: `pyro.infer.abstract_infer.TracePosterior`

Parameters

- **model** – probabilistic model defined as a function
- **guide** – guide used for sampling defined as a function
- **num_samples** – number of samples to draw from the guide (default 10)

This method performs posterior inference by importance sampling using the guide as the proposal distribution. If no guide is provided, it defaults to proposing from the model’s prior.

4.4 Inference Utilities

class `EmpiricalMarginal` (*trace_posterior*, *sites=None*, *validate_args=None*)

Bases: `pyro.distributions.empirical.Empirical`

Marginal distribution, that wraps over a `TracePosterior` object to provide a a marginal over one or more latent sites or the return values of the `TracePosterior`’s model. If multiple sites are specified, they must have the same tensor shape.

Parameters

- **trace_posterior** (`TracePosterior`) – a `TracePosterior` instance representing a Monte Carlo posterior.
- **sites** (*list*) – optional list of sites for which we need to generate the marginal distribution. Note that for multiple sites, the shape for the site values must match (needed by the underlying `Empirical` class).

class `TracePosterior`

Bases: `object`

Abstract `TracePosterior` object from which posterior inference algorithms inherit. When run, collects a bag of execution traces from the approximate posterior. This is designed to be used by other utility classes like `EmpiricalMarginal`, that need access to the collected execution traces.

run (**args*, ***kwargs*)

Calls `self._traces` to populate execution traces from a stochastic Pyro model.

Parameters

- **args** – optional args taken by `self._traces`.
- **kwargs** – optional keywords args taken by `self._traces`.

class `TracePredictive` (*model*, *posterior*, *num_samples*)

Bases: `pyro.infer.abstract_infer.TracePosterior`

Generates and holds traces from the posterior predictive distribution, given model execution traces from the approximate posterior. This is achieved by constraining latent sites to randomly sampled parameter values from

the model execution traces and running the model forward to generate traces with new response (“_RETURN”) sites.

Parameters

- **model** – arbitrary Python callable containing Pyro primitives.
- **posterior** (`TracePosterior`) – trace posterior instance holding samples from the model’s approximate posterior.
- **num_samples** (`int`) – number of samples to generate.

4.5 MCMC

4.5.1 MCMC

class `MCMC` (*kernel, num_samples, warmup_steps=0*)

Bases: `pyro.infer.abstract_infer.TracePosterior`

Wrapper class for Markov Chain Monte Carlo algorithms. Specific MCMC algorithms are `TraceKernel` instances and need to be supplied as a `kernel` argument to the constructor.

Parameters

- **kernel** – An instance of the `TraceKernel` class, which when given an execution trace returns another sample trace from the target (posterior) distribution.
- **num_samples** (`int`) – The number of samples that need to be generated, excluding the samples discarded during the warmup phase.
- **warmup_steps** (`int`) – Number of warmup iterations. The samples generated during the warmup phase are discarded.

4.5.2 HMC

class `HMC` (*model, step_size=None, trajectory_length=None, num_steps=None, adapt_step_size=False, transforms=None*)

Bases: `pyro.infer.mcmc.trace_kernel.TraceKernel`

Simple Hamiltonian Monte Carlo kernel, where `step_size` and `num_steps` need to be explicitly specified by the user.

References

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

Parameters

- **model** – Python callable containing pyro primitives.
- **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.
- **trajectory_length** (`float`) – Length of a MCMC trajectory. If not specified, it will be set to `step_size` x `num_steps`. In case `num_steps` is not specified, it will be set to 2π .

- **num_steps** (*int*) – The number of discrete steps over which to simulate Hamiltonian dynamics. The state at the end of the trajectory is returned as the proposal. This value is always equal to `int(trajectory_length / step_size)`.
- **adapt_step_size** (*bool*) – A flag to decide if we want to adapt `step_size` during warm-up phase using Dual Averaging scheme.
- **transforms** (*dict*) – Optional dictionary that specifies a transform for a sample site with constrained support to unconstrained space. The transform should be invertible, and implement `log_abs_det_jacobian`. If not specified and the model has sites with constrained support, automatic transformations will be applied, as specified in `torch.distributions.constraint_registry`.

Example:

```
true_coefs = torch.tensor([1., 2., 3.])
data = torch.randn(2000, 3)
dim = 3
labels = dist.Bernoulli(logits=(true_coefs * data).sum(-1)).sample()

def model(data):
    coefs_mean = torch.zeros(dim)
    coefs = pyro.sample('beta', dist.Normal(coefs_mean, torch.ones(3)))
    y = pyro.sample('y', dist.Bernoulli(logits=(coefs * data).sum(-1)),
↳obs=labels)
    return y

hmc_kernel = HMC(model, step_size=0.0855, num_steps=4)
mcmc_run = MCMC(hmc_kernel, num_samples=500, warmup_steps=100).run(data)
posterior = EmpiricalMarginal(mcmc_run, 'beta')
print(posterior.mean)
```

cleanup ()

Optional method to clean up any residual state on termination.

diagnostics ()

Relevant diagnostics (optional) to be printed at regular intervals of the MCMC run. Returns *None* by default.

Returns String containing the diagnostic summary. e.g. acceptance rate

Return type string

end_warmup ()

Optional method to tell kernel that warm-up phase has been finished.

initial_trace ()

Returns an initial trace from the prior to initiate the MCMC run.

Returns Trace instance.

sample (*trace*)

Samples a trace from the approximate posterior distribution, when given an existing trace.

Parameters

- **trace** – Current execution trace.
- **time_step** (*int*) – Current time step.

Returns New trace sampled from the approximate posterior distribution.

`setup (*args, **kwargs)`

Optional method to set up any state required at the start of the simulation run.

Parameters

- ***args** – Algorithm specific positional arguments.
- ****kwargs** – Algorithm specific keyword arguments.

4.5.3 NUTS

class `NUTS` (*model*, *step_size=None*, *adapt_step_size=False*, *transforms=None*)

Bases: `pyro.infer.mcmc.hmc.HMC`

No-U-Turn Sampler kernel, which provides an efficient and convenient way to run Hamiltonian Monte Carlo. The number of steps taken by the integrator is dynamically adjusted on each call to `sample` to ensure an optimal length for the Hamiltonian trajectory [1]. As such, the samples generated will typically have lower autocorrelation than those generated by the `HMC` kernel. Optionally, the NUTS kernel also provides the ability to adapt step size during the warmup phase.

Refer to the [baseball example](#) to see how to do Bayesian inference in Pyro using NUTS.

References

[1] *The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo*, Matthew D. Hoffman, and Andrew Gelman

Parameters

- **model** – Python callable containing pyro primitives.
- **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.
- **transforms** (*dict*) – Optional dictionary that specifies a transform for a sample site with constrained support to unconstrained space. The transform should be invertible, and implement `log_abs_det_jacobian`. If not specified and the model has sites with constrained support, automatic transformations will be applied, as specified in `torch.distributions.constraint_registry`.

Example:

```

true_coefs = torch.tensor([1., 2., 3.])
data = torch.randn(2000, 3)
dim = 3
labels = dist.Bernoulli(logits=(true_coefs * data).sum(-1)).sample()

def model(data):
    coefs_mean = torch.zeros(dim)
    coefs = pyro.sample('beta', dist.Normal(coefs_mean, torch.ones(3)))
    y = pyro.sample('y', dist.Bernoulli(logits=(coefs * data).sum(-1)),
↳obs=labels)
    return y

nuts_kernel = NUTS(model, adapt_step_size=True)
mcmc_run = MCMC(nuts_kernel, num_samples=500, warmup_steps=300).run(data)

```

(continues on next page)

(continued from previous page)

```
posterior = EmpiricalMarginal(mcmc_run, 'beta')
print(posterior.mean)
```

sample (*trace*)

Samples a trace from the approximate posterior distribution, when given an existing trace.

Parameters

- **trace** – Current execution trace.
- **time_step** (*int*) – Current time step.

Returns New trace sampled from the approximate posterior distribution.

5.1 PyTorch Distributions

Most distributions in Pyro are thin wrappers around PyTorch distributions. For details on the PyTorch distribution interface, see `torch.distributions.distribution.Distribution`. For differences between the Pyro and PyTorch interfaces, see `TorchDistributionMixin`.

5.1.1 Bernoulli

```
class Bernoulli (probs=None, logits=None, validate_args=None)  
    Wraps torch.distributions.bernoulli.Bernoulli with TorchDistributionMixin.
```

5.1.2 Beta

```
class Beta (concentration1, concentration0, validate_args=None)  
    Wraps torch.distributions.beta.Beta with TorchDistributionMixin.
```

5.1.3 Categorical

```
class Categorical (probs=None, logits=None, validate_args=None)  
    Wraps torch.distributions.categorical.Categorical with TorchDistributionMixin.
```

5.1.4 Cauchy

```
class Cauchy (loc, scale, validate_args=None)  
    Wraps torch.distributions.cauchy.Cauchy with TorchDistributionMixin.
```

5.1.5 Chi2

```
class Chi2 (df, validate_args=None)  
    Wraps torch.distributions.chi2.Chi2 with TorchDistributionMixin.
```

5.1.6 Dirichlet

```
class Dirichlet (concentration, validate_args=None)  
    Wraps torch.distributions.dirichlet.Dirichlet with TorchDistributionMixin.
```

5.1.7 Exponential

```
class Exponential (rate, validate_args=None)  
    Wraps torch.distributions.exponential.Exponential with TorchDistributionMixin.
```

5.1.8 ExponentialFamily

```
class ExponentialFamily (batch_shape=torch.Size([]), event_shape=torch.Size([]), vali-  
                        date_args=None)  
    Wraps torch.distributions.exp_family.ExponentialFamily with  
    TorchDistributionMixin.
```

5.1.9 FisherSnedecor

```
class FisherSnedecor (df1, df2, validate_args=None)  
    Wraps torch.distributions.fishersnedecor.FisherSnedecor with  
    TorchDistributionMixin.
```

5.1.10 Gamma

```
class Gamma (concentration, rate, validate_args=None)  
    Wraps torch.distributions.gamma.Gamma with TorchDistributionMixin.
```

5.1.11 Geometric

```
class Geometric (probs=None, logits=None, validate_args=None)  
    Wraps torch.distributions.geometric.Geometric with TorchDistributionMixin.
```

5.1.12 Gumbel

```
class Gumbel (loc, scale, validate_args=None)  
    Wraps torch.distributions.gumbel.Gumbel with TorchDistributionMixin.
```

5.1.13 Independent

```
class Independent (base_distribution, reinterpreted_batch_ndims, validate_args=None)  
    Wraps torch.distributions.independent.Independent with TorchDistributionMixin.
```

5.1.14 Laplace

```
class Laplace (loc, scale, validate_args=None)
    Wraps torch.distributions.laplace.Laplace with TorchDistributionMixin.
```

5.1.15 LogNormal

```
class LogNormal (loc, scale, validate_args=None)
    Wraps torch.distributions.log_normal.LogNormal with TorchDistributionMixin.
```

5.1.16 LogisticNormal

```
class LogisticNormal (loc, scale, validate_args=None)
    Wraps torch.distributions.logistic_normal.LogisticNormal with
    TorchDistributionMixin.
```

5.1.17 Multinomial

```
class Multinomial (total_count=1, probs=None, logits=None, validate_args=None)
    Wraps torch.distributions.multinomial.Multinomial with TorchDistributionMixin.
```

5.1.18 MultivariateNormal

```
class MultivariateNormal (loc, covariance_matrix=None, precision_matrix=None, scale_tril=None,
    validate_args=None)
    Wraps torch.distributions.multivariate_normal.MultivariateNormal with
    TorchDistributionMixin.
```

5.1.19 Normal

```
class Normal (loc, scale, validate_args=None)
    Wraps torch.distributions.normal.Normal with TorchDistributionMixin.
```

5.1.20 OneHotCategorical

```
class OneHotCategorical (probs=None, logits=None, validate_args=None)
    Wraps torch.distributions.one_hot_categorical.OneHotCategorical with
    TorchDistributionMixin.
```

5.1.21 Pareto

```
class Pareto (scale, alpha, validate_args=None)
    Wraps torch.distributions.pareto.Pareto with TorchDistributionMixin.
```

5.1.22 Poisson

```
class Poisson (rate, validate_args=None)  
    Wraps torch.distributions.poisson.Poisson with TorchDistributionMixin.
```

5.1.23 RelaxedBernoulli

```
class RelaxedBernoulli (temperature, probs=None, logits=None, validate_args=None)  
    Wraps torch.distributions.relaxed_bernoulli.RelaxedBernoulli with  
    TorchDistributionMixin.
```

5.1.24 RelaxedOneHotCategorical

```
class RelaxedOneHotCategorical (temperature, probs=None, logits=None, validate_args=None)  
    Wraps torch.distributions.relaxed_categorical.RelaxedOneHotCategorical with  
    TorchDistributionMixin.
```

5.1.25 StudentT

```
class StudentT (df, loc=0.0, scale=1.0, validate_args=None)  
    Wraps torch.distributions.studentT.StudentT with TorchDistributionMixin.
```

5.1.26 TransformedDistribution

```
class TransformedDistribution (base_distribution, transforms, validate_args=None)  
    Wraps torch.distributions.transformed_distribution.TransformedDistribution  
    with TorchDistributionMixin.
```

5.1.27 Uniform

```
class Uniform (low, high, validate_args=None)  
    Wraps torch.distributions.uniform.Uniform with TorchDistributionMixin.
```

5.2 Pyro Distributions

5.2.1 Abstract Distribution

```
class Distribution
```

```
    Bases: object
```

```
    Base class for parameterized probability distributions.
```

```
    Distributions in Pyro are stochastic function objects with sample() and log_prob() methods. Distribution  
    are stochastic functions with fixed parameters:
```

```
d = dist.Bernoulli(param)  
x = d()                               # Draws a random sample.  
p = d.log_prob(x)                     # Evaluates log probability of x.
```


Implementing New Distributions:

Derived classes must implement the methods: `sample()`, `log_prob()`.

Examples:

Take a look at the [examples](#) to see how they interact with inference algorithms.

`__call__` (**args*, ***kwargs*)

Samples a random value (just an alias for `.sample(*args, **kwargs)`).

For tensor distributions, the returned tensor should have the same `.shape` as the parameters.

Returns A random value.

Return type `torch.Tensor`

`enumerate_support` ()

Returns a representation of the parametrized distribution’s support, along the first dimension. This is implemented only by discrete distributions.

Note that this returns support values of all the batched RVs in lock-step, rather than the full cartesian product.

Returns An iterator over the distribution’s discrete support.

Return type iterator

`has_enumerate_support` = **False**

`has_rsample` = **False**

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

Evaluates log probability densities for each of a batch of samples.

Parameters ***x*** (`torch.Tensor`) – A single value or a batch of values batched along axis 0.

Returns log probability densities as a one-dimensional `Tensor` with same batch size as value and params. The shape of the result should be `self.batch_size`.

Return type `torch.Tensor`

`sample` (**args*, ***kwargs*)

Samples a random value.

For tensor distributions, the returned tensor should have the same `.shape` as the parameters, unless otherwise noted.

Parameters **`sample_shape`** (`torch.Size`) – the size of the iid batch to be drawn from the distribution.

Returns A random value or batch of random values (if parameters are batched). The shape of the result should be `self.shape()`.

Return type `torch.Tensor`

`score_parts` (*x*, **args*, ***kwargs*)

Computes ingredients for stochastic gradient estimators of ELBO.

The default implementation is correct both for non-reparameterized and for fully reparameterized distributions. Partially reparameterized distributions should override this method to compute correct `.score_function` and `.entropy_term` parts.

Parameters ***x*** (`torch.Tensor`) – A single value or batch of values.

Returns A `ScoreParts` object containing parts of the ELBO estimator.

Return type ScoreParts

5.2.2 TorchDistribution

class TorchDistribution (*batch_shape=torch.Size([])*, *event_shape=torch.Size([])*, *validate_args=None*)

Bases: `torch.distributions.distribution.Distribution`, `pyro.distributions.torch_distribution.TorchDistributionMixin`

Base class for PyTorch-compatible distributions with Pyro support.

This should be the base class for almost all new Pyro distributions.

Note: Parameters and data should be of type `Tensor` and all methods return type `Tensor` unless otherwise noted.

Tensor Shapes:

TorchDistributions provide a method `.shape()` for the tensor shape of samples:

```
x = d.sample(sample_shape)
assert x.shape == d.shape(sample_shape)
```

Pyro follows the same distribution shape semantics as PyTorch. It distinguishes between three different roles for tensor shapes of samples:

- *sample shape* corresponds to the shape of the iid samples drawn from the distribution. This is taken as an argument by the distribution's *sample* method.
- *batch shape* corresponds to non-identical (independent) parameterizations of the distribution, inferred from the distribution's parameter shapes. This is fixed for a distribution instance.
- *event shape* corresponds to the event dimensions of the distribution, which is fixed for a distribution class. These are collapsed when we try to score a sample from the distribution via *d.log_prob(x)*.

These shapes are related by the equation:

```
assert d.shape(sample_shape) == sample_shape + d.batch_shape + d.event_shape
```

Distributions provide a vectorized `meth'~torch.distributions.distribution.Distribution.log_prob'` method that evaluates the log probability density of each event in a batch independently, returning a tensor of shape `sample_shape + d.batch_shape`:

```
x = d.sample(sample_shape)
assert x.shape == d.shape(sample_shape)
log_p = d.log_prob(x)
assert log_p.shape == sample_shape + d.batch_shape
```

Implementing New Distributions:

Derived classes must implement the methods `sample()` (or `rsample()` if `.has_rsample == True`) and `log_prob()`, and must implement the properties `batch_shape`, and `event_shape`. Discrete classes may also implement the `enumerate_support()` method to improve gradient estimates and set `.has_enumerate_support = True`.

5.2.3 Binomial

class Binomial (*total_count=1, probs=None, logits=None, validate_args=None*)

Bases: `torch.distributions.distribution.Distribution`, `pyro.distributions.torch_distribution.TorchDistributionMixin`

Creates a Binomial distribution parameterized by *total_count* and either *probs* or *logits* (but not both). *total_count* must be broadcastable with *probs/logits*.

This is adapted from `torch.distributions.binomial.Binomial`, with the important difference that *total_count* is not limited to being a single *int*, but can be a *torch.Tensor*.

Example:

```
>>> m = Binomial(100, torch.Tensor([0, .2, .8, 1]))
>>> x = m.sample()
0
22
71
100
[torch.FloatTensor of size 4]

>>> m = Binomial(torch.Tensor([[5.], [10.]]), torch.Tensor([0.5, 0.8]))
>>> x = m.sample()
4 5
7 6
[torch.FloatTensor of size (2,2)]
```

Parameters

- **total_count** (*Tensor*) – number of Bernoulli trials
- **probs** (*Tensor*) – Event probabilities
- **logits** (*Tensor*) – Event log-odds

arg_constraints = {'probs': <torch.distributions.constraints.Interval object at 0x7f...

enumerate_support ()

Returns tensor containing all values supported by a discrete distribution. The result will enumerate over dimension 0, so the shape of the result will be (*cardinality*,) + *batch_shape* + *event_shape* (where *event_shape* = () for univariate distributions).

Note that this enumerates over all batched tensors in lock-step $[[0, 0], [1, 1], \dots]$. To iterate over the full Cartesian product use `itertools.product(m.enumerate_support())`.

Returns: Tensor iterating over dimension 0.

has_enumerate_support = True

log_prob (*value*)

Returns the log of the probability density/mass function evaluated at *value*.

Args: value (Tensor):

logits

mean

param_shape

probs

sample (*sample_shape=torch.Size([])*)

Generates a *sample_shape* shaped sample or *sample_shape* shaped batch of samples if the distribution parameters are batched.

support

variance

5.2.4 Delta

class Delta (*v, log_density=0.0, event_dim=0, validate_args=None*)

Bases: `pyro.distributions.torch_distribution.TorchDistribution`

Degenerate discrete distribution (a single point).

Discrete distribution that assigns probability one to the single element in its support. Delta distribution parameterized by a random choice should not be used with MCMC based inference, as doing so produces incorrect results.

Parameters

- **v** (*torch.Tensor*) – The single support element.
- **log_density** (*torch.Tensor*) – An optional density for this Delta. This is useful to keep the class of *Delta* distributions closed under differentiable transformation.
- **event_dim** (*int*) – Optional event dimension, defaults to zero.

arg_constraints = {'log_density': <torch.distributions.constraints._Real object at 0x...

has_rsample = True

log_prob (*x*)

Returns the log of the probability density/mass function evaluated at *value*.

Args: *value* (Tensor):

mean

rsample (*sample_shape=torch.Size([])*)

Generates a *sample_shape* shaped reparameterized sample or *sample_shape* shaped batch of reparameterized samples if the distribution parameters are batched.

support = <torch.distributions.constraints._Real object>

variance

5.2.5 EmpiricalDistribution

class Empirical (*validate_args=None*)

Bases: `pyro.distributions.torch_distribution.TorchDistribution`

Empirical distribution associated with the sampled data.

add (*value, weight=None, log_weight=None*)

Adds a new data point to the sample. The values in successive calls to `add` must have the same tensor shape and size. Optionally, an importance weight can be specified via `log_weight` or `weight` (default value of `1` is used if not specified).

Parameters

- **value** (*torch.Tensor*) – tensor to add to the sample.

- **weight** (*torch.Tensor*) – log weight (optional) corresponding to the sample.
- **log_weight** (*torch.Tensor*) – weight (optional) corresponding to the sample.

arg_constraints = {}

enumerate_support ()

See `pyro.distributions.torch_distribution.TorchDistribution.enumerate_support()`

event_shape

See `pyro.distributions.torch_distribution.TorchDistribution.event_shape()`

get_samples_and_weights ()

has_enumerate_support = True

log_prob (*value*)

Returns the log of the probability mass function evaluated at *value*. Note that this currently only supports scoring values with empty *sample_shape*, i.e. an arbitrary batched sample is not allowed.

Parameters *value* (*torch.Tensor*) – scalar or tensor value to be scored.

mean

See `pyro.distributions.torch_distribution.TorchDistribution.mean()`

sample (*sample_shape=torch.Size([])*)

See `pyro.distributions.torch_distribution.TorchDistribution.sample()`

sample_size

Number of samples that constitute the empirical distribution.

Return int number of samples collected.

support = <torch.distributions.constraints._Real object>

variance

See `pyro.distributions.torch_distribution.TorchDistribution.variance()`

5.2.6 HalfCauchy

class HalfCauchy (*loc, scale*)

Bases: `pyro.distributions.torch.TransformedDistribution`

Half-Cauchy distribution.

This is a continuous distribution with lower-bounded domain ($x > loc$). See also the Cauchy distribution.

Parameters

- **loc** (*torch.Tensor*) – lower bound of the distribution.
- **scale** (*torch.Tensor*) – half width at half maximum.

arg_constraints = {'loc': <torch.distributions.constraints._Real object at 0x7f7b2bc4>

entropy ()

Returns entropy of distribution, batched over *batch_shape*.

Returns: Tensor of shape *batch_shape*.

loc

log_prob (*value*)

Scores the sample by inverting the transform(s) and computing the score using the score of the base distribution and the log abs det jacobian.

scale

support

5.2.7 LowRankMultivariateNormal

class LowRankMultivariateNormal (*loc, W_term, D_term, trace_term=None*)

Bases: `pyro.distributions.torch_distribution.TorchDistribution`

Low Rank Multivariate Normal distribution.

Implements fast computation for log probability of Multivariate Normal distribution when the covariance matrix has the form:

```
covariance_matrix = W.T @ W + D.
```

Here D is a diagonal vector and W is a matrix of size $M \times N$. The computation will be beneficial when $M \ll N$.

Parameters

- **loc** (*torch.Tensor*) – Mean. Must be a 1D or 2D tensor with the last dimension of size N .
- **W_term** (*torch.Tensor*) – W term of covariance matrix. Must be in 2 dimensional of size $M \times N$.
- **D_term** (*torch.Tensor*) – D term of covariance matrix. Must be in 1 dimensional of size N .
- **trace_term** (*float*) – A optional term to be added into Mahalabonis term according to $p(y) = N(y|loc, cov).exp(-1/2 * trace_term)$.

arg_constraints = {'covariance_matrix_D_term': `<torch.distributions.constraints._Grea`

has_rsample = `True`

log_prob (*value*)

Returns the log of the probability density/mass function evaluated at *value*.

Args: value (Tensor):

mean

rsample (*sample_shape=torch.Size([])*)

Generates a *sample_shape* shaped reparameterized sample or *sample_shape* shaped batch of reparameterized samples if the distribution parameters are batched.

scale_tril

support = `<torch.distributions.constraints._Real object>`

variance

5.2.8 OMTMultivariateNormal

class OMTMultivariateNormal (*loc, scale_tril*)

Bases: `pyro.distributions.torch.MultivariateNormal`

Multivariate normal (Gaussian) distribution with OMT gradients w.r.t. both parameters. Note the gradient computation w.r.t. the Cholesky factor has cost $O(D^3)$, although the resulting gradient variance is generally expected to be lower.

A distribution over vectors in which all the elements have a joint Gaussian density.

Parameters

- **loc** (*torch.Tensor*) – Mean.
- **scale_tril** (*torch.Tensor*) – Cholesky of Covariance matrix.

arg_constraints = {'loc': <torch.distributions.constraints._Real object at 0x7f7b2bc4...>

rsample (*sample_shape=torch.Size([])*)

Generates a *sample_shape* shaped reparameterized sample or *sample_shape* shaped batch of reparameterized samples if the distribution parameters are batched.

5.2.9 Rejector

class Rejector (*propose, log_prob_accept, log_scale*)

Bases: `pyro.distributions.torch_distribution.TorchDistribution`

Rejection sampled distribution given an acceptance rate function.

Parameters

- **propose** (*Distribution*) – A proposal distribution that samples batched proposals via `propose()`. `rsample()` supports a *sample_shape* arg only if `propose()` supports a *sample_shape* arg.
- **log_prob_accept** (*callable*) – A callable that inputs a batch of proposals and returns a batch of log acceptance probabilities.
- **log_scale** – Total log probability of acceptance.

has_rsample = True

log_prob (*x*)

Returns the log of the probability density/mass function evaluated at *value*.

Args: *value* (Tensor):

rsample (*sample_shape=torch.Size([])*)

Generates a *sample_shape* shaped reparameterized sample or *sample_shape* shaped batch of reparameterized samples if the distribution parameters are batched.

score_parts (*x*)

Computes ingredients for stochastic gradient estimators of ELBO.

The default implementation is correct both for non-reparameterized and for fully reparameterized distributions. Partially reparameterized distributions should override this method to compute correct `.score_function` and `.entropy_term` parts.

Parameters **x** (*torch.Tensor*) – A single value or batch of values.

Returns A `ScoreParts` object containing parts of the ELBO estimator.

Return type ScoreParts

5.2.10 VonMises

class VonMises (*loc, concentration, validate_args=None*)

Bases: `pyro.distributions.torch_distribution.TorchDistribution`

A circular von Mises distribution.

Currently only `log_prob()` is implemented.

Parameters

- **loc** (`torch.Tensor`) – an angle in radians.
- **concentration** (`torch.Tensor`) – concentration parameter

arg_constraints = {'concentration': `<torch.distributions.constraints._GreaterThan obj`

log_prob (*value*)

Returns the log of the probability density/mass function evaluated at *value*.

Args: value (Tensor):

support = `<torch.distributions.constraints._Real object>`

5.3 Transformed Distributions

5.3.1 InverseAutoregressiveFlow

class InverseAutoregressiveFlow (*input_dim, hidden_dim, sigmoid_bias=2.0, permutation=None*)

Bases: `torch.distributions.transforms.Transform`

An implementation of an Inverse Autoregressive Flow. Together with the `TransformedDistribution` this provides a way to create richer variational approximations.

Example usage:

```
>>> base_dist = Normal(...)
>>> iaf = InverseAutoregressiveFlow(...)
>>> pyro.module("my_iaf", iaf.module)
>>> iaf_dist = TransformedDistribution(base_dist, [iaf])
```

Note that this implementation is only meant to be used in settings where the inverse of the Bijector is never explicitly computed (rather the result is cached from the forward call). In the context of variational inference, this means that the `InverseAutoregressiveFlow` should only be used in the guide, i.e. in the variational distribution. In other contexts the inverse could in principle be computed but this would be a (potentially) costly computation that scales with the dimension of the input (and in any case support for this is not included in this implementation).

Parameters

- **input_dim** (`int`) – dimension of input
- **hidden_dim** (`int`) – hidden dimension (number of hidden units)
- **sigmoid_bias** (`float`) – bias on the hidden units fed into the sigmoid; default='2.0'

- **permutation** (*bool*) – whether the order of the inputs should be permuted (by default the conditional dependence structure of the autoregression follows the sequential order)

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
2. Variational Inference with Normalizing Flows [arXiv:1505.05770] Danilo Jimenez Rezende, Shakir Mohamed
3. MADE: Masked Autoencoder for Distribution Estimation [arXiv:1502.03509] Mathieu Germain, Karol Gregor, Iain Murray, Hugo Larochelle

arn

Return type pyro.nn.AutoRegressiveNN

Return the AutoRegressiveNN associated with the InverseAutoregressiveFlow

log_abs_det_jacobian (*x*, *y*)

Calculates the elementwise determinant of the log jacobian

Parameters in Pyro are basically thin wrappers around PyTorch Variables that carry unique names. As such Parameters are the primary stateful objects in Pyro. Users typically interact with parameters via the Pyro primitive *pyro.param*. Parameters play a central role in stochastic variational inference, where they are used to represent point estimates for the parameters in parameterized families of models and guides.

6.1 ParamStore

class ParamStoreDict

Bases: `object`

Global store for parameters in Pyro. This is basically a key-value store. The typical user interacts with the ParamStore primarily through the primitive *pyro.param*.

See [Intro Part II](#) for further discussion and [SVI Part I](#) for some examples.

Some things to bear in mind when using parameters in Pyro:

- parameters must be assigned unique names
- the *init_tensor* argument to *pyro.param* is only used the first time that a given (named) parameter is registered with Pyro.
- for this reason, a user may need to use the *clear()* method if working in a REPL in order to get the desired behavior. this method can also be invoked with *pyro.clear_param_store()*.
- the internal name of a parameter within a PyTorch *nn.Module* that has been registered with Pyro is prepended with the Pyro name of the module. so nothing prevents the user from having two different modules each of which contains a parameter named *weight*. by contrast, a user can only have one top-level parameter named *weight* (outside of any module).
- parameters can be saved and loaded from disk using *save* and *load*.

clear()

Clear the ParamStore

get_all_param_names ()

Get all parameter names in the ParamStore

get_param (*name*, *init_tensor=None*, *constraint=<torch.distributions.constraints._Real object>*)

Get parameter from its name. If it does not yet exist in the ParamStore, it will be created and stored. The Pyro primitive *pyro.param* dispatches to this method.

Parameters

- **name** (*str*) – parameter name
- **init_tensor** (*torch.Tensor*) – initial tensor

Returns parameter

Return type `torch.Tensor`

get_state ()

Get the ParamStore state.

load (*filename*)

Loads parameters from disk

Parameters **filename** – file name to load from

named_parameters ()

Returns an iterator over tuples of the form (name, parameter) for each parameter in the ParamStore

param_name (*p*)

Get parameter name from parameter

Parameters **p** – parameter

Returns parameter name

replace_param (*param_name*, *new_param*, *old_param*)

Replace the param *param_name* with current value *old_param* with the new value *new_param*

Parameters

- **param_name** (*str*) – parameter name
- **new_param** (*torch.Tensor*) – the parameter to be put into the ParamStore
- **old_param** – the parameter to be removed from the ParamStore

save (*filename*)

Save parameters to disk

Parameters **filename** – file name to save to

set_state (*state*)

Set the ParamStore state using state from a previous *get_state()* call

module_from_param_with_module_name (*param_name*)

param_with_module_name (*pyro_name*, *param_name*)

user_param_name (*param_name*)

The module *pyro.nn* provides implementations of neural network modules that are useful in the context of deep probabilistic programming. None of these modules is really part of the core language.

7.1 AutoRegressiveNN

```
class AutoRegressiveNN(input_dim, hidden_dim, output_dim_multiplier=1, mask_encoding=None,  
                       permutation=None)
```

Bases: `torch.nn.modules.module.Module`

A simple implementation of a MADE-like auto-regressive neural network.

Reference: MADE: Masked Autoencoder for Distribution Estimation [arXiv:1502.03509] Mathieu Germain, Karol Gregor, Iain Murray, Hugo Larochelle

Parameters

- **input_dim** (*int*) – the dimensionality of the input
- **hidden_dim** (*int*) – the dimensionality of the hidden units
- **output_dim_multiplier** (*int*) – the dimensionality of the output is given by `input_dim x output_dim_multiplier`. specifically the shape of the output for a single vector input is `[output_dim_multiplier, input_dim]`. for any `i, j` in `range(0, output_dim_multiplier)` the subset of outputs `[i, :]` has identical autoregressive structure to `[j, :]`. defaults to `1`
- **mask_encoding** (*torch.LongTensor*) – a torch Tensor that controls the autoregressive structure (see reference). by default this is chosen at random.
- **permutation** (*torch.LongTensor*) – an optional permutation that is applied to the inputs and controls the order of the autoregressive factorization. in particular for the identity permutation the autoregressive structure is such that the Jacobian is upper triangular. by default this is chosen at random.

forward (*z*)
the forward method

`get_mask_encoding()`

Get the mask encoding associated with the neural network: basically the quantity $m(k)$ in the MADE paper.

`get_permutation()`

Get the permutation applied to the inputs (by default this is chosen at random)

class `MaskedLinear` (*in_features*, *out_features*, *mask*, *bias=True*)

Bases: `torch.nn.modules.linear.Linear`

A linear mapping with a given mask on the weights (arbitrary bias)

Parameters

- **in_features** (*int*) – the number of input features
- **out_features** (*int*) – the number of output features
- **mask** (*torch.Tensor*) – the mask to apply to the *in_features* x *out_features* weight matrix
- **bias** (*bool*) – whether or not *MaskedLinear* should include a bias term. defaults to *True*

forward (*_input*)

the forward method that does the masked linear computation and returns the result

The module `pyro.optim` provides support for optimization in Pyro. In particular it provides `PyroOptim`, which is used to wrap PyTorch optimizers and manage optimizers for dynamically generated parameters (see the tutorial [SVI Part I](#) for a discussion). Any custom optimization algorithms are also to be found here.

8.1 Pyro Optimizers

AdagradRMSProp (*optim_args*)

A wrapper for an optimizer that is a mash-up of `Adagrad` and `RMSprop`.

ClippedAdam (*optim_args*)

A wrapper for a modification of the `Adam` optimization algorithm that supports gradient clipping.

class PyroOptim (*optim_constructor*, *optim_args*)

Bases: `object`

A wrapper for `torch.optim.Optimizer` objects that helps with managing dynamically generated parameters.

Parameters

- **optim_constructor** – a `torch.optim.Optimizer`
- **optim_args** – a dictionary of learning arguments for the optimizer or a callable that returns such dictionaries

get_state ()

Get state associated with all the optimizers in the form of a dictionary with key-value pairs (parameter name, optim state dicts)

load (*filename*)

Parameters *filename* – file name to load from

Load optimizer state from disk

save (*filename*)

Parameters *filename* – file name to save to

Save optimizer state to disk

set_state (*state_dict*)

Set the state associated with all the optimizers using the state obtained from a previous call to `get_state()`

8.2 Pytorch Optimizers

LBFGS (*optim_args*)

Wraps `torch.optim.LBFGS` with *PyroOptim*.

Adamax (*optim_args*)

Wraps `torch.optim.Adamax` with *PyroOptim*.

Adagrad (*optim_args*)

Wraps `torch.optim.Adagrad` with *PyroOptim*.

SGD (*optim_args*)

Wraps `torch.optim.SGD` with *PyroOptim*.

Adam (*optim_args*)

Wraps `torch.optim.Adam` with *PyroOptim*.

Rprop (*optim_args*)

Wraps `torch.optim.Rprop` with *PyroOptim*.

ASGD (*optim_args*)

Wraps `torch.optim.ASGD` with *PyroOptim*.

RMSprop (*optim_args*)

Wraps `torch.optim.RMSprop` with *PyroOptim*.

SparseAdam (*optim_args*)

Wraps `torch.optim.SparseAdam` with *PyroOptim*.

Adadelta (*optim_args*)

Wraps `torch.optim.Adadelta` with *PyroOptim*.

Poutine (Effect handlers)

Beneath the built-in inference algorithms, Pyro has a library of composable effect handlers for creating new inference algorithms and working with probabilistic programs. Pyro's inference algorithms are all built by applying these handlers to stochastic functions.

9.1 Handlers

Poutine is a library of composable effect handlers for recording and modifying the behavior of Pyro programs. These lower-level ingredients simplify the implementation of new inference algorithms and behavior.

Handlers can be used as higher-order functions, decorators, or context managers to modify the behavior of functions or blocks of code:

For example, consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

We can mark sample sites as observed using `condition`, which returns a callable with the same input and output signatures as `model`:

```
>>> conditioned_model = poutine.condition(model, data={"z": 1.0})
```

We can also use handlers as decorators:

```
>>> @pyro.condition(data={"z": 1.0})
... def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

Or as context managers:

```
>>> with pyro.condition(data={"z": 1.0}):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     y = z ** 2
```

Handlers compose freely:

```
>>> conditioned_model = poutine.condition(model, data={"z": 1.0})
>>> traced_model = poutine.trace(conditioned_model)
```

Many inference algorithms or algorithmic components can be implemented in just a few lines of code:

```
>>> guide_tr = poutine.trace(guide).get_trace(...)
>>> model_tr = poutine.trace(poutine.replay(conditioned_model, trace=tr)).get_trace(..
↪.)
>>> monte_carlo_elbo = model_tr.log_prob_sum() - guide_tr.log_prob_sum()
```

block (*fn=None, hide=None, expose=None, hide_types=None, expose_types=None*)

This handler selectively hides pyro primitive sites from the outside world. Default behavior: block everything

A site is hidden if at least one of the following holds:

1. `msg["name"]` in `hide`
2. `msg["type"]` in `hide_types`
3. `msg["name"]` not in `expose` and `msg["type"]` not in `expose_types`
4. `hide`, `hide_types`, and `expose_types` are all `None`

For example, suppose the stochastic function `fn` has two sample sites “a” and “b”. Then any effect outside of `BlockMessenger(fn, hide=["a"])` will not be applied to site “a” and will only see site “b”:

```
>>> fn_inner = trace(fn)
>>> fn_outer = trace(block(fn_inner, hide=["a"]))
>>> trace_inner = fn_inner.get_trace()
>>> trace_outer = fn_outer.get_trace()
>>> "a" in trace_inner
True
>>> "a" in trace_outer
False
>>> "b" in trace_inner
True
>>> "b" in trace_outer
True
```

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **hide** – list of site names to hide
- **expose** – list of site names to be exposed while all others hidden
- **hide_types** – list of site types to be hidden
- **expose_types** – list of site types to be exposed while all others hidden

Returns stochastic function decorated with a *BlockMessenger*

condition (*fn=None, data=None*)

Given a stochastic function with some sample statements and a dictionary of observations at names, change the sample statements at those names into observes with those values

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

To observe a value for site *z*, we can write

```
>>> conditioned_model = condition(model, data={"z": value})
```

This is equivalent to adding *obs=value* as a keyword argument to *pyro.sample("z", ...)* in *model*.

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **data** – a dict or a *Trace*

Returns stochastic function decorated with a *ConditionMessenger*

do (*fn=None, data=None*)

Given a stochastic function with some sample statements and a dictionary of values at names, set the return values of those sites equal to the values and hide them from the rest of the stack as if they were hard-coded to those values by using *block*.

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

To intervene with a value for site *z*, we can write

```
>>> intervened_model = do(model, data={"z": value})
```

This is equivalent to replacing *z = pyro.sample("z", ...)* with *z = value*.

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **data** – a dict or a *Trace*

Returns stochastic function decorated with a *BlockMessenger* and *pyro.poutine.condition_messenger.ConditionMessenger*

enum (*fn=None, first_available_dim=None*)

Enumerates in parallel over discrete sample sites marked *infer={"enumerate": "parallel"}*.

Parameters **first_available_dim** (*int*) – The first tensor dimension (counting from the right) that is available for parallel enumeration. This dimension and all dimensions left may be used internally by Pyro.

escape (*fn=None, escape_fn=None*)

Given a callable that contains Pyro primitive calls, evaluate *escape_fn* on each site, and if the result is *True*, raise a *NonlocalExit* exception that stops execution and returns the offending site.

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **escape_fn** – function that takes a partial trace and a site, and returns a boolean value to decide whether to exit at that site

Returns stochastic function decorated with *EscapeMessenger*

indep (*fn=None, name=None, size=None, dim=None*)

Note: Low-level; use *iarange* instead.

This messenger keeps track of stack of independence information declared by nested *irange* and *iarange* contexts. This information is stored in a *cond_indep_stack* at each sample/observe site for consumption by *TraceMessenger*.

infer_config (*fn=None, config_fn=None*)

Given a callable that contains Pyro primitive calls and a callable taking a trace site and returning a dictionary, updates the value of the infer kwarg at a sample site to *config_fn(site)*

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **config_fn** – a callable taking a site and returning an infer dict

Returns stochastic function decorated with *InferConfigMessenger*

lift (*fn=None, prior=None*)

Given a stochastic function with param calls and a prior distribution, create a stochastic function where all param calls are replaced by sampling from prior. Prior should be a callable or a dict of names to callables.

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
>>> lifted_model = lift(model, prior={"s": dist.Exponential(0.3)})
```

lift makes param statements behave like sample statements using the distributions in *prior*. In this example, site *s* will now behave as if it was replaced with `s = pyro.sample("s", dist.Exponential(0.3))`:

```
>>> tr = trace(lifted_model).get_trace(0.0)
>>> tr.nodes["s"]["type"] == "sample"
True
>>> tr2 = trace(lifted_model).get_trace(0.0)
>>> (tr2.nodes["s"]["value"] == tr.nodes["s"]["value"]).all()
False
```

Parameters

- **fn** – function whose parameters will be lifted to random values
- **prior** – prior function in the form of a Distribution or a dict of stochastic fns

Returns *fn* decorated with a *LiftMessenger*

replay (*fn=None, trace=None, sites=None*)

Given a callable that contains Pyro primitive calls, return a callable that runs the original, reusing the values at sites in trace at those sites in the new trace

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

replay makes sample statements behave as if they had sampled the values at the corresponding sites in the trace:

```
>>> replayed_model = replay(model, trace=old_trace)
>>> replayed_model(0.0) == old_trace.nodes["_RETURN"]["value"]
True
```

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **trace** – a *Trace* data structure to replay against
- **sites** – list or dict of names of sample sites in fn to replay against, defaulting to all sites

Returns a stochastic function decorated with a *ReplayMessenger*

queue (*fn=None, queue=None, max_tries=None, extend_fn=None, escape_fn=None, num_samples=None*)

Used in sequential enumeration over discrete variables.

Given a stochastic function and a queue, return a return value from a complete trace in the queue

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **queue** – a queue data structure like multiprocessing.Queue to hold partial traces
- **max_tries** – maximum number of attempts to compute a single complete trace
- **extend_fn** – function (possibly stochastic) that takes a partial trace and a site, and returns a list of extended traces
- **escape_fn** – function (possibly stochastic) that takes a partial trace and a site, and returns a boolean value to decide whether to exit
- **num_samples** – optional number of extended traces for extend_fn to return

Returns stochastic function decorated with poutine logic

scale (*fn=None, scale=None*)

Given a stochastic function with some sample statements and a positive scale factor, scale the score of all sample and observe sites in the function.

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s), obs=1.0)
...     return z ** 2
```

scale multiplicatively scales the log-probabilities of sample sites:

```
>>> scaled_model = scale(model, scale=0.5)
>>> scaled_tr = trace(scaled_model).get_trace(0.0)
>>> unscaled_tr = trace(model).get_trace(0.0)
>>> (scaled_tr.log_prob_sum() == 0.5 * unscaled_tr.log_prob_sum()).all()
True
```

Parameters

- **fn** – a stochastic function (callable containing Pyro primitive calls)
- **scale** – a positive scaling factor

Returns stochastic function decorated with a *ScaleMessenger*

trace (*fn=None, graph_type=None, param_only=None*)

Return a handler that records the inputs and outputs of primitive calls and their dependencies.

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

We can record its execution using `trace` and use the resulting data structure to compute the log-joint probability of all of the sample sites in the execution or extract all parameters.

```
>>> trace = pyro.poutine.trace(model).get_trace(0.0)
>>> logp = trace.log_prob_sum()
>>> params = [trace.nodes[name]["value"].unconstrained() for name in trace.param_
↳nodes]
```

Parameters

- **fn** – a stochastic function (callable containing pyro primitive calls)
- **graph_type** – string that specifies the kind of graph to construct
- **param_only** – if true, only records params and not samples

Returns stochastic function decorated with a *TraceMessenger*

9.2 Trace

class Trace (**args, **kwargs*)

Bases: `object`

Execution trace data structure built on top of `networkx.DiGraph`.

An execution trace of a Pyro program is a record of every call to `pyro.sample()` and `pyro.param()` in a single execution of that program. Traces are directed graphs whose nodes represent primitive calls or input/output, and whose edges represent conditional dependence relationships between those primitive calls. They are created and populated by `poutine.trace`.

Each node (or site) in a trace contains the name, input and output value of the site, as well as additional metadata added by inference algorithms or user annotation. In the case of `pyro.sample`, the trace also includes the stochastic function at the site, and any observed data added by users.

Consider the following Pyro program:

```
>>> def model(x):
...     s = pyro.param("s", torch.tensor(0.5))
...     z = pyro.sample("z", dist.Normal(x, s))
...     return z ** 2
```

We can record its execution using `pyro.poutine.trace` and use the resulting data structure to compute the log-joint probability of all of the sample sites in the execution or extract all parameters.

```
>>> trace = pyro.poutine.trace(model).get_trace(0.0)
>>> logp = trace.log_prob_sum()
>>> params = [trace.nodes[name]["value"].unconstrained() for name in trace.param_
↳nodes]
```

We can also inspect or manipulate individual nodes in the trace. `trace.nodes` contains a `collections.OrderedDict` of site names and metadata corresponding to `x`, `s`, `z`, and the return value:

```
>>> print(list(name for name in trace.nodes.keys()))
["_INPUT", "s", "z", "_RETURN"]
```

As in `networkx.DiGraph`, values of `trace.nodes` are dictionaries of node metadata:

```
>>> print(trace.nodes["z"])
{'type': 'sample', 'name': 'z', 'is_observed': False,
 'fn': Normal(), 'value': tensor(0.6480), 'args': (), 'kwargs': {},
 'infer': {}, 'scale': 1.0, 'cond_indep_stack': (),
 'done': True, 'stop': False, 'continuation': None}
```

'infer' is a dictionary of user- or algorithm-specified metadata. 'args' and 'kwargs' are the arguments passed via `pyro.sample` to `fn.__call__` or `fn.log_prob`. 'scale' is used to scale the log-probability of the site when computing the log-joint. 'cond_indep_stack' contains data structures corresponding to `pyro.iarange` contexts appearing in the execution. 'done', 'stop', and 'continuation' are only used by Pyro's internals.

add_edge

Identical to `networkx.DiGraph.remove_edge()`

add_node(*site_name*, **args*, ***kwargs*)

Parameters *site_name* (*string*) – the name of the site to be added

Adds a site to the trace.

Identical to `networkx.DiGraph.add_node()` but raises an error when attempting to add a duplicate node instead of silently overwriting.

compute_log_prob(*site_filter*=<function <lambda>>)

Compute the site-wise log probabilities of the trace. Each `log_prob` has shape equal to the corresponding `batch_shape`. Each `log_prob_sum` is a scalar. Both computations are memoized.

compute_score_parts()

Compute the batched local score parts at each site of the trace. Each `log_prob` has shape equal to the corresponding `batch_shape`. Each `log_prob_sum` is a scalar. All computations are memoized.

copy()

Makes a shallow copy of self with nodes and edges preserved. Identical to `networkx.DiGraph.copy()`, but preserves the type and the `self.graph_type` attribute

edges

Identical to `networkx.DiGraph.edges`

graph

Identical to `networkx.DiGraph.graph`

in_degree

Identical to `networkx.DiGraph.in_degree()`

is_directed

Identical to `networkx.DiGraph.is_directed`

iter_stochastic_nodes()

Returns an iterator over stochastic nodes in the trace.

log_prob_sum (*site_filter=<function <lambda>>*)

Compute the site-wise log probabilities of the trace. Each `log_prob` has shape equal to the corresponding `batch_shape`. Each `log_prob_sum` is a scalar. The computation of `log_prob_sum` is memoized.

Returns total log probability.

Return type `torch.Tensor`

nodes

Identical to `networkx.DiGraph.nodes`

nonreparam_stochastic_nodes

Returns a list of names of sample sites whose stochastic functions are not reparameterizable primitive distributions

observation_nodes

Returns a list of names of observe sites

param_nodes

Returns a list of names of param sites

remove_node

Identical to `networkx.DiGraph.remove_node()`

reparameterized_nodes

Returns a list of names of sample sites whose stochastic functions are reparameterizable primitive distributions

stochastic_nodes

Returns a list of names of sample sites

successors

Identical to `networkx.DiGraph.successors()`

9.3 Messengers

Messenger objects contain the implementations of the effects exposed by handlers. Advanced users may modify the implementations of messengers behind existing handlers or write new messengers that implement new effects and compose correctly with the rest of the library.

9.3.1 Messenger

class Handler (*msngr, fn*)

Bases: `object`

Wrapper class created by `Messenger.__call__`

class Messenger

Bases: `object`

Context manager class that modifies behavior and adds side effects to stochastic functions i.e. callables containing pyro primitive statements.

This is the base Messenger class. It implements the default behavior for all pyro primitives, so that the joint distribution induced by a stochastic function `fn` is identical to the joint distribution induced by `Messenger()(fn)`.

Class of transformers for messages passed during inference. Most inference operations are implemented in subclasses of this.

9.3.2 BlockMessenger

class BlockMessenger (*hide_all=True, expose_all=False, hide=None, expose=None, hide_types=None, expose_types=None*)

Bases: `pyro.poutine.messenger.Messenger`

This Messenger selectively hides pyro primitive sites from the outside world. Default behavior: block everything. `BlockMessenger` has a flexible interface that allows users to specify in several different ways which sites should be hidden or exposed.

A site is hidden if at least one of the following holds:

1. `msg["name"]` in `hide`
2. `msg["type"]` in `hide_types`
3. `msg["name"]` not in `expose` and `msg["type"]` not in `expose_types`
4. `hide_all == True` and `hide`, `hide_types`, and `expose_types` are all `None`

For example, suppose the stochastic function `fn` has two sample sites “a” and “b”. Then any poutine outside of `BlockMessenger(fn, hide=["a"])` will not be applied to site “a” and will only see site “b”:

```
>>> fn_inner = TraceMessenger()(fn)
>>> fn_outer = TraceMessenger()(BlockMessenger(hide=["a"])(TraceMessenger()(fn)))
>>> trace_inner = fn_inner.get_trace()
>>> trace_outer = fn_outer.get_trace()
>>> "a" in trace_inner
True
>>> "a" in trace_outer
False
>>> "b" in trace_inner
True
>>> "b" in trace_outer
True
```

See the constructor for details.

Parameters

- **hide_all** (*bool*) – hide all sites

- **expose_all** (*bool*) – expose all sites normally
- **hide** (*list*) – list of site names to hide, rest will be exposed normally
- **expose** (*list*) – list of site names to expose, rest will be hidden
- **hide_types** (*list*) – list of site types to hide, rest will be exposed normally
- **expose_types** (*list*) – list of site types to expose normally, rest will be hidden

9.3.3 ConditionMessenger

class ConditionMessenger (*data*)

Bases: *pyro.poutine.messenger.Messenger*

Adds values at observe sites to condition on data and override sampling

9.3.4 EscapeMessenger

class EscapeMessenger (*escape_fn*)

Bases: *pyro.poutine.messenger.Messenger*

Messenger that does a nonlocal exit by raising a `util.NonlocalExit` exception

9.3.5 IndepMessenger

class CondIndepStackFrame

Bases: *pyro.poutine.indep_messenger.CondIndepStackFrame*

vectorized

class IndepMessenger (*name, size, dim=None*)

Bases: *pyro.poutine.messenger.Messenger*

This messenger keeps track of stack of independence information declared by nested `irange` and `iarange` contexts. This information is stored in a `cond_indep_stack` at each sample/observe site for consumption by `TraceMessenger`.

next_context ()

Increments the counter.

9.3.6 LiftMessenger

class LiftMessenger (*prior*)

Bases: *pyro.poutine.messenger.Messenger*

Messenger which “lifts” parameters to random samples. Given a stochastic function with param calls and a prior, creates a stochastic function where all param calls are replaced by sampling from prior.

Prior should be a callable or a dict of names to callables.

9.3.7 ReplayMessenger

class `ReplayMessenger` (*trace=None, sites=None*)

Bases: `pyro.poutine.messenger.Messenger`

Messenger for replaying from an existing execution trace.

9.3.8 ScaleMessenger

class `ScaleMessenger` (*scale*)

Bases: `pyro.poutine.messenger.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` (*float or torch.Tensor*) – a positive scaling factor

9.3.9 TraceMessenger

class `TraceHandler` (*msngr, fn*)

Bases: `pyro.poutine.messenger.Handler`

Execution trace poutine.

A `TraceHandler` records the input and output to every pyro primitive and stores them as a site in a `Trace()`. This should, in theory, be sufficient information for every inference algorithm (along with the implicit computational graph in the `Variables?`)

We can also use this for visualization.

get_trace (**args, **kwargs*)

Returns data structure

Return type `pyro.poutine.Trace`

Helper method for a very common use case. Calls this poutine and returns its trace instead of the function's return value.

trace

class `TraceMessenger` (*graph_type=None, param_only=None*)

Bases: `pyro.poutine.messenger.Messenger`

Execution trace messenger.

A `TraceMessenger` records the input and output to every pyro primitive and stores them as a site in a `Trace()`. This should, in theory, be sufficient information for every inference algorithm (along with the implicit computational graph in the `Variables?`)

We can also use this for visualization.

get_trace ()

Returns data structure

Return type `pyro.poutine.Trace`

Helper method for a very common use case. Returns a shallow copy of `self.trace`.

identify_dense_edges (*trace*)

Modifies a trace in-place by adding all edges based on the *cond_indep_stack* information stored at each site.

9.4 Runtime

exception NonlocalExit (*site, *args, **kwargs*)

Bases: `exceptions.Exception`

Exception for exiting nonlocally from poutine execution.

Used by `poutine.EscapeMessenger` to return site information.

reset_stack ()

Reset the state of the frames remaining in the stack. Necessary for multiple re-executions in `poutine.queue`.

am_i_wrapped ()

Checks whether the current computation is wrapped in a poutine. :returns: bool

apply_stack (*initial_msg*)

Execute the effect stack at a single site according to the following scheme:

1. For each `Messenger` in the stack from bottom to top, execute `Messenger._process_message` with the message; if the message field “stop” is True, stop; otherwise, continue
2. Apply default behavior (`default_process_message`) to finish remaining site execution
3. For each `Messenger` in the stack from top to bottom, execute `_postprocess_message` to update the message and internal messenger state with the site results
4. If the message field “continuation” is not None, call it with the message

Parameters *initial_msg* (*dict*) – the starting version of the trace site

Returns None

default_process_message (*msg*)

Default method for processing messages in inference. :param msg: a message to be processed :returns: None

validate_message (*msg*)

Asserts that the message has a valid format. :returns: None

9.5 Utilities

all_escape (*trace, msg*)

Parameters

- **trace** – a partial trace
- **msg** – the message at a pyro primitive site

Returns boolean decision value

Utility function that checks if a site is not already in a trace.

Used by `EscapeMessenger` to decide whether to do a nonlocal exit at a site. Subroutine for approximately integrating out variables for variance reduction.

discrete_escape (*trace, msg*)

Parameters

- **trace** – a partial trace
- **msg** – the message at a pyro primitive site

Returns boolean decision value

Utility function that checks if a sample site is discrete and not already in a trace.

Used by `EscapeMessenger` to decide whether to do a nonlocal exit at a site. Subroutine for integrating out discrete variables for variance reduction.

enum_extend(*trace, msg, num_samples=None*)

Parameters

- **trace** – a partial trace
- **msg** – the message at a pyro primitive site
- **num_samples** – maximum number of extended traces to return.

Returns a list of traces, copies of input trace with one extra site

Utility function to copy and extend a trace with sites based on the input site whose values are enumerated from the support of the input site's distribution.

Used for exact inference and integrating out discrete variables.

mc_extend(*trace, msg, num_samples=None*)

Parameters

- **trace** – a partial trace
- **msg** – the message at a pyro primitive site
- **num_samples** – maximum number of extended traces to return.

Returns a list of traces, copies of input trace with one extra site

Utility function to copy and extend a trace with sites based on the input site whose values are sampled from the input site's function.

Used for Monte Carlo marginalization of individual sample sites.

prune_subsample_sites(*trace*)

Copies and removes all subsample sites from a trace.

site_is_subsample(*site*)

Determines whether a trace site originated from a subsample statement inside an *iarange*.

Warning: Code in `pyro.contrib` is under various stages of development. This code makes no guarantee about maintaining backwards compatibility.

10.1 Automatic Guide Generation

The `pyro.contrib.autoguide` module provides algorithms to automatically generate guides from simple models, for use in *SVI*. For example to generate a mean field Gaussian guide:

```
def model():
    ...

guide = AutoDiagonalNormal(model) # a mean field guide
svi = SVI(model, guide, Adam({'lr': 1e-3}), Trace_ELBO())
```

Automatic guides can also be combined using `pyro.poutine.block()` and `AutoGuideList`.

10.1.1 AutoGuide

class `AutoGuide` (*model*)

Bases: `object`

Base class for automatic guides.

Derived classes must implement the `__call__()` method.

Auto guides can be used individually or combined in an `AutoGuideList` object.

`__call__` (**args*, ***kwargs*)

A guide with the same **args*, ***kwargs* as the base model.

Returns A dict mapping sample site name to sampled value.

Return type `dict`

sample_latent (***kwargs*)

Samples an encoded latent given the same **args*, ***kwargs* as the base model.

10.1.2 AutoGuideList

class `AutoGuideList` (*model*)

Bases: `pyro.contrib.autoguide.AutoGuide`

Container class to combine multiple automatic guides.

Example usage:

```
guide = AutoGuideList(my_model)
guide.add(AutoDiagonalNormal(poutine.block(model, hide=["assignment"])))
guide.add(AutoDiscreteParallel(poutine.block(model, expose=["assignment"])))
svi = SVI(model, guide, optim, Trace_ELBO())
```

Parameters `model` (*callable*) – a Pyro model

__call__ (**args*, ***kwargs*)

A composite guide with the same **args*, ***kwargs* as the base model.

Returns A dict mapping sample site name to sampled value.

Return type `dict`

add (*part*)

Add an automatic guide for part of the model. The guide should have been created by blocking the model to restrict to a subset of sample sites. No two parts should operate on any one sample site.

Parameters `part` (`AutoGuide`) – a partial guide to add

10.1.3 AutoContinuous

class `AutoContinuous` (*model*)

Bases: `pyro.contrib.autoguide.AutoGuide`

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

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

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

Parameters `model` (*callable*) – a Pyro model

Reference:

[1] ‘Automatic Differentiation Variational Inference’, Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, David M. Blei

__call__ (**args*, ***kwargs*)

An automatic guide with the same **args*, ***kwargs* as the base model.

Returns A dict mapping sample site name to sampled value.

Return type `dict`

sample_latent (**args*, ***kwargs*)

Samples an encoded latent given the same **args*, ***kwargs* as the base model.

10.1.4 AutoMultivariateNormal

class `AutoMultivariateNormal` (*model*)

Bases: `pyro.contrib.autoguide.AutoContinuous`

This implementation of `AutoContinuous` uses a Cholesky factorization of a Multivariate Normal 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, ...)
```

By default the mean vector is initialized to zero and the Cholesky factor is initialized to the identity. To change this default behavior the user should call `pyro.param()` before beginning inference, e.g.:

```
latent_dim = 10
pyro.param("auto_loc", torch.randn(latent_dim))
pyro.param("auto_scale_tril", torch.tril(torch.rand(latent_dim)),
          constraint=constraints.lower_cholesky)
```

median (**args*, ***kwargs*)

Returns the posterior median value of each latent variable.

Returns A dict mapping sample site name to median tensor.

Return type `dict`

quantiles (*quantiles*, **args*, ***kwargs*)

Returns posterior quantiles each latent variable. Example:

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

Parameters `quantiles` (*torch.Tensor* or *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`

sample_latent (**args*, ***kwargs*)

Samples the (single) multivariate normal latent used in the auto guide.

10.1.5 AutoDiagonalNormal

class `AutoDiagonalNormal` (*model*)

Bases: `pyro.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, ...)
```

By default the mean vector is initialized to zero and the scale is initialized to the identity. To change this default behavior the user should call `pyro.param()` before beginning inference, e.g.:

```
latent_dim = 10
pyro.param("auto_loc", torch.randn(latent_dim))
pyro.param("auto_scale", torch.ones(latent_dim),
           constraint=constraints.positive)
```

median (*args, **kwargs)

Returns the posterior median value of each latent variable.

Returns A dict mapping sample site name to median tensor.

Return type dict

quantiles (quantiles, *args, **kwargs)

Returns posterior quantiles each latent variable. Example:

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

Parameters **quantiles** (*torch.Tensor* or *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

sample_latent (*args, **kwargs)

Samples the (single) diagonal normal latent used in the auto guide.

10.1.6 AutoDiscreteParallel

class **AutoDiscreteParallel** (*model*)

Bases: *pyro.contrib.autoguide.AutoGuide*

A discrete mean-field guide that learns a latent discrete distribution for each discrete site in the model.

__call__ (*args, **kwargs)

An automatic guide with the same *args, **kwargs as the base model.

Returns A dict mapping sample site name to sampled value.

Return type dict

10.2 Gaussian Processes

See the [Gaussian Processes tutorial](#) for an introduction.

10.2.1 Models

class **GPMModel** (*X, y, kernel, mean_function=None, jitter=1e-06, name=None*)

Bases: *pyro.contrib.gp.util.Parameterized*

Base class for Gaussian Process models.

The core of a Gaussian Process is a covariance function k which governs the similarity between input points. Given k , we can establish a distribution over functions f by a multivariate normal distribution

$$p(f(X)) = \mathcal{N}(0, k(X, X)),$$

where X is any set of input points and $k(X, X)$ is a covariance matrix whose entries are outputs $k(x, z)$ of k over input pairs (x, z) . This distribution is usually denoted by

$$f \sim \mathcal{GP}(0, k).$$

Note: Generally, beside a covariance matrix k , a Gaussian Process can also be specified by a mean function m (which is a zero-value function by default). In that case, its distribution will be

$$p(f(X)) = \mathcal{N}(m(X), k(X, X)).$$

Gaussian Process models are *Parameterized* subclasses. So its parameters can be learned, set priors, or fixed by using corresponding methods from *Parameterized*. A typical way to define a Gaussian Process model is

```
>>> X = torch.tensor([[1., 5, 3], [4, 3, 7]])
>>> y = torch.tensor([2., 1])
>>> kernel = gp.kernels.RBF(input_dim=3)
>>> gpr = gp.models.GPRegression(X, y, kernel)
```

There are two ways to train a Gaussian Process model:

- Using an MCMC algorithm (in module `pyro.infer.mcmc`) on `model()` to get posterior samples for the Gaussian Process's parameters. For example:

```
>>> hmc_kernel = HMC(gpr.model)
>>> mcmc_run = MCMC(hmc_kernel, num_samples=10)
>>> posterior_ls_trace = [] # store lengthscale trace
>>> ls_name = param_with_module_name(gpr.kernel.name, "lengthscale")
>>> for trace, _ in mcmc_run._traces():
...     posterior_ls_trace.append(trace.nodes[ls_name]["value"])
```

- Using a variational inference (e.g. *SVI*) on the pair `model()`, `guide()` as in *SVI tutorial*:

```
>>> optimizer = pyro.optim.Adam({"lr": 0.01})
>>> svi = SVI(gpr.model, gpr.guide, optimizer, loss=Trace_ELBO())
>>> for i in range(1000):
...     svi.step()
```

To give a prediction on new dataset, simply use `forward()` like any PyTorch `torch.nn.Module`:

```
>>> Xnew = torch.tensor([[2., 3, 1]])
>>> f_loc, f_cov = gpr(Xnew, full_cov=True)
```

Reference:

- [1] *Gaussian Processes for Machine Learning*, Carl E. Rasmussen, Christopher K. I. Williams

Parameters

- **X** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.
- **kernel** (*Kernel*) – A Pyro kernel object, which is the covariance function k .

- **mean_function** (*callable*) – An optional mean function m of this Gaussian process. By default, we use zero mean.
- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.
- **name** (*str*) – Name of this model.

model ()

A “model” stochastic function. If `self.y` is `None`, this method returns mean and variance of the Gaussian Process prior.

guide ()

A “guide” stochastic function to be used in variational inference methods. It also gives posterior information to the method `forward()` for prediction.

forward (*Xnew*, *full_cov=False*)

Computes the mean and covariance matrix (or variance) of Gaussian Process posterior on a test input data X_{new} :

$$p(f^* | X_{new}, X, y, k, \theta),$$

where θ are parameters of this model.

Note: Model’s parameters θ together with kernel’s parameters have been learned from a training procedure (MCMC or SVI).

Parameters

- **Xnew** (*torch.Tensor*) – A input data for testing. Note that `Xnew.shape[1:]` must be the same as `X.shape[1:]`.
- **full_cov** (*bool*) – A flag to decide if we want to predict full covariance matrix or just variance.

Returns loc and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

set_data (*X*, *y=None*)

Sets data for Gaussian Process models.

Some examples to utilize this method are:

- Batch training on a sparse variational model:

```

>>> Xu = torch.tensor([[1., 0, 2]]) # inducing input
>>> likelihood = gp.likelihoods.Gaussian()
>>> vsgp = gp.models.SparseVariationalGP(X, y, kernel, Xu, likelihood)
>>> svi = SVI(vsgp.model, vsgp.guide, optimizer, Trace_ELBO())
>>> batched_X, batched_y = X.split(split_size=10), y.split(split_size=10)
>>> for Xi, yi in zip(batched_X, batched_y):
...     vsgp.set_data(Xi, yi)
...     svi.step()

```

- Making a two-layer Gaussian Process stochastic function:

```

>>> gpr1 = gp.models.GPRegression(X, None, kernel1, name="GPR1")
>>> Z, _ = gp_layer1.model()
>>> gpr2 = gp.models.GPRegression(Z, y, kernel2, name="GPR2")
>>> def two_layer_model():
...     Z, _ = gpr1.model()
...     gpr2.set_data(Z, y)
...     return gpr2.model()

```

References:

[1] *Scalable Variational Gaussian Process Classification*, James Hensman, Alexander G. de G. Matthews, Zoubin Ghahramani

[2] *Deep Gaussian Processes*, Andreas C. Damianou, Neil D. Lawrence

Parameters

- **X** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.

optimize (*optimizer=None, loss=None, num_steps=1000*)

A convenient method to optimize parameters for the Gaussian Process model using SVI.

Parameters

- **optimizer** (*PyroOptim*) – A Pyro optimizer.
- **loss** (*ELBO*) – A Pyro loss instance.
- **num_steps** (*int*) – Number of steps to run SVI.

Returns a list of losses during the training procedure

Return type *list*

class GPRegression (*X, y, kernel, noise=None, mean_function=None, jitter=1e-06, name='GPR'*)

Bases: *pyro.contrib.gp.models.model.GPModel*

Gaussian Process Regression model.

The core of a Gaussian Process is a covariance function k which governs the similarity between input points. Given k , we can establish a distribution over functions f by a multivariate normal distribution

$$p(f(X)) = \mathcal{N}(0, k(X, X)),$$

where X is any set of input points and $k(X, X)$ is a covariance matrix whose entries are outputs $k(x, z)$ of k over input pairs (x, z) . This distribution is usually denoted by

$$f \sim \mathcal{GP}(0, k).$$

Note: Generally, beside a covariance matrix k , a Gaussian Process can also be specified by a mean function m (which is a zero-value function by default). In that case, its distribution will be

$$p(f(X)) = \mathcal{N}(m(X), k(X, X)).$$

Given inputs X and their noisy observations y , the Gaussian Process Regression model takes the form

$$\begin{aligned}f &\sim \mathcal{GP}(0, k(X, X)), \\y &\sim f + \epsilon,\end{aligned}$$

where ϵ is Gaussian noise.

Note: This model has $\mathcal{O}(N^3)$ complexity for training, $\mathcal{O}(N^3)$ complexity for testing. Here, N is the number of train inputs.

Reference:

[1] *Gaussian Processes for Machine Learning*, Carl E. Rasmussen, Christopher K. I. Williams

Parameters

- **X** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.
- **kernel** (*Kernel*) – A Pyro kernel object, which is the covariance function k .
- **noise** (*torch.Tensor*) – Variance of Gaussian noise of this model.
- **mean_function** (*callable*) – An optional mean function m of this Gaussian process. By default, we use zero mean.
- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.
- **name** (*str*) – Name of this model.

model()

A “model” stochastic function. If `self.y` is `None`, this method returns mean and variance of the Gaussian Process prior.

guide()

A “guide” stochastic function to be used in variational inference methods. It also gives posterior information to the method `forward()` for prediction.

forward(Xnew, full_cov=False, noiseless=True)

Computes the mean and covariance matrix (or variance) of Gaussian Process posterior on a test input data X_{new} :

$$p(f^* | X_{new}, X, y, k, \epsilon) = \mathcal{N}(loc, cov).$$

Note: The noise parameter `noise` (ϵ) together with kernel’s parameters have been learned from a training procedure (MCMC or SVI).

Parameters

- **Xnew** (*torch.Tensor*) – A input data for testing. Note that `Xnew.shape[1:]` must be the same as `self.X.shape[1:]`.
- **full_cov** (*bool*) – A flag to decide if we want to predict full covariance matrix or just variance.

- **noiseless** (*bool*) – A flag to decide if we want to include noise in the prediction output or not.

Returns loc and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

class SparseGPRegression (*X, y, kernel, Xu, noise=None, mean_function=None, approx=None, jitter=1e-06, name='SGPR'*)
 Bases: `pyro.contrib.gp.models.model.GPModel`

Sparse Gaussian Process Regression model.

In *GPRegression* model, when the number of input data X is large, the covariance matrix $k(X, X)$ will require a lot of computational steps to compute its inverse (for log likelihood and for prediction). By introducing an additional inducing-input parameter X_u , we can reduce computational cost by approximate $k(X, X)$ by a low-rank Nymström approximation Q (see reference [1]), where

$$Q = k(X, X_u)k(X, X)^{-1}k(X_u, X).$$

Given inputs X , their noisy observations y , and the inducing-input parameters X_u , the model takes the form:

$$\begin{aligned} u &\sim \mathcal{GP}(0, k(X_u, X_u)), \\ f &\sim q(f | X, X_u) = \mathbb{E}_{p(u)}q(f | X, X_u, u), \\ y &\sim f + \epsilon, \end{aligned}$$

where ϵ is Gaussian noise and the conditional distribution $q(f | X, X_u, u)$ is an approximation of

$$p(f | X, X_u, u) = \mathcal{N}(m, k(X, X) - Q),$$

whose terms m and $k(X, X) - Q$ is derived from the joint multivariate normal distribution:

$$[f, u] \sim \mathcal{GP}(0, k([X, X_u], [X, X_u])).$$

This class implements three approximation methods:

- Deterministic Training Conditional (DTC):

$$q(f | X, X_u, u) = \mathcal{N}(m, 0),$$

which in turns will imply

$$f \sim \mathcal{N}(0, Q).$$

- Fully Independent Training Conditional (FITC):

$$q(f | X, X_u, u) = \mathcal{N}(m, \text{diag}(k(X, X) - Q)),$$

which in turns will correct the diagonal part of the approximation in DTC:

$$f \sim \mathcal{N}(0, Q + \text{diag}(k(X, X) - Q)).$$

- Variational Free Energy (VFE), which is similar to DTC but has an additional *trace_term* in the model’s log likelihood. This additional term makes “VFE” equivalent to the variational approach in `SparseVariationalGP` (see reference [2]).

Note: This model has $\mathcal{O}(NM^2)$ complexity for training, $\mathcal{O}(NM^2)$ complexity for testing. Here, N is the number of train inputs, M is the number of inducing inputs.

References:

[1] *A Unifying View of Sparse Approximate Gaussian Process Regression*, Joaquin Quiñero-Candela, Carl E. Rasmussen

[2] *Variational learning of inducing variables in sparse Gaussian processes*, Michalis Titsias

Parameters

- **x** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.
- **kernel** (*Kernel*) – A Pyro kernel object, which is the covariance function k .
- **Xu** (*torch.Tensor*) – Initial values for inducing points, which are parameters of our model.
- **noise** (*torch.Tensor*) – Variance of Gaussian noise of this model.
- **mean_function** (*callable*) – An optional mean function m of this Gaussian process. By default, we use zero mean.
- **approx** (*str*) – One of approximation methods: “DTC”, “FITC”, and “VFE” (default).
- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.
- **name** (*str*) – Name of this model.

model ()

A “model” stochastic function. If `self.y` is `None`, this method returns mean and variance of the Gaussian Process prior.

guide ()

A “guide” stochastic function to be used in variational inference methods. It also gives posterior information to the method `forward()` for prediction.

forward (*Xnew*, *full_cov=False*, *noiseless=True*)

Computes the mean and covariance matrix (or variance) of Gaussian Process posterior on a test input data X_{new} :

$$p(f^* | X_{new}, X, y, k, X_u, \epsilon) = \mathcal{N}(loc, cov).$$

Note: The noise parameter `noise` (ϵ), the inducing-point parameter `Xu`, together with `kernel`’s parameters have been learned from a training procedure (MCMC or SVI).

Parameters

- **Xnew** (*torch.Tensor*) – A input data for testing. Note that `Xnew.shape[1:]` must be the same as `self.X.shape[1:]`.
- **full_cov** (*bool*) – A flag to decide if we want to predict full covariance matrix or just variance.
- **noiseless** (*bool*) – A flag to decide if we want to include noise in the prediction output or not.

Returns loc and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

```
class VariationalGP(X, y, kernel, likelihood, mean_function=None, latent_shape=None,
                    whiten=False, jitter=1e-06, name='VGP')
```

Bases: `pyro.contrib.gp.models.model.GPModel`

Variational Gaussian Process model.

This model deals with both Gaussian and non-Gaussian likelihoods. Given inputs X and their noisy observations y , the model takes the form

$$\begin{aligned} f &\sim \mathcal{GP}(0, k(X, X)), \\ y &\sim p(y) = p(y | f)p(f), \end{aligned}$$

where $p(y | f)$ is the likelihood.

We will use a variational approach in this model by approximating $q(f)$ to the posterior $p(f | y)$. Precisely, $q(f)$ will be a multivariate normal distribution with two parameters `f_loc` and `f_scale_tril`, which will be learned during a variational inference process.

Note: This model can be seen as a special version of `SparseVariationalGP` model with $X_u = X$.

Note: This model has $\mathcal{O}(N^3)$ complexity for training, $\mathcal{O}(N^3)$ complexity for testing. Here, N is the number of train inputs. Size of variational parameters is $\mathcal{O}(N^2)$.

Parameters

- **X** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.
- **kernel** (*Kernel*) – A Pyro kernel object, which is the covariance function k .
- **Likelihood likelihood** (*likelihood*) – A likelihood object.
- **mean_function** (*callable*) – An optional mean function m of this Gaussian process. By default, we use zero mean.
- **latent_shape** (*torch.Size*) – Shape for latent processes (*batch_shape* of $q(f)$). By default, it equals to output batch shape `y.shape[:-1]`. For the multi-class classification problems, `latent_shape[-1]` should correspond to the number of classes.
- **whiten** (*bool*) – A flag to tell if variational parameters `f_loc` and `f_scale_tril` are transformed by the inverse of `Lff`, where `Lff` is the lower triangular decomposition of $kernel(X, X)$. Enable this flag will help optimization.

- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.
- **name** (*str*) – Name of this model.

model ()

A “model” stochastic function. If `self.y` is `None`, this method returns mean and variance of the Gaussian Process prior.

guide ()

A “guide” stochastic function to be used in variational inference methods. It also gives posterior information to the method `forward()` for prediction.

forward (*Xnew*, *full_cov=False*)

Computes the mean and covariance matrix (or variance) of Gaussian Process posterior on a test input data X_{new} :

$$p(f^* | X_{new}, X, y, k, f_{loc}, f_{scale_tril}) = \mathcal{N}(loc, cov).$$

Note: Variational parameters `f_loc`, `f_scale_tril`, together with kernel’s parameters have been learned from a training procedure (MCMC or SVI).

Parameters

- **Xnew** (*torch.Tensor*) – A input data for testing. Note that `Xnew.shape[1:]` must be the same as `self.X.shape[1:]`.
- **full_cov** (*bool*) – A flag to decide if we want to predict full covariance matrix or just variance.

Returns `loc` and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

class VariationalSparseGP (*X*, *y*, *kernel*, *Xu*, *likelihood*, *mean_function=None*, *latent_shape=None*, *num_data=None*, *whiten=False*, *jitter=1e-06*, *name='SVGP'*)

Bases: `pyro.contrib.gp.models.model.GPModel`

Variational Sparse Gaussian Process model.

In *VariationalGP* model, when the number of input data X is large, the covariance matrix $k(X, X)$ will require a lot of computational steps to compute its inverse (for log likelihood and for prediction). This model introduces an additional inducing-input parameter X_u to solve that problem. Given inputs X , their noisy observations y , and the inducing-input parameters X_u , the model takes the form:

$$\begin{aligned} [f, u] &\sim \mathcal{GP}(0, k([X, X_u], [X, X_u])), \\ y &\sim p(y) = p(y | f)p(f), \end{aligned}$$

where $p(y | f)$ is the likelihood.

We will use a variational approach in this model by approximating $q(f, u)$ to the posterior $p(f, u | y)$. Precisely, $q(f) = p(f | u)q(u)$, where $q(u)$ is a multivariate normal distribution with two parameters `u_loc` and `u_scale_tril`, which will be learned during a variational inference process.

Note: This model can be learned using MCMC method as in reference [2]. See also `GPModel`.

Note: This model has $\mathcal{O}(NM^2)$ complexity for training, $\mathcal{O}(M^3)$ complexity for testing. Here, N is the number of train inputs, M is the number of inducing inputs. Size of variational parameters is $\mathcal{O}(M^2)$.

References:

[1] *Scalable variational Gaussian process classification*, James Hensman, Alexander G. de G. Matthews, Zoubin Ghahramani

[2] *MCMC for Variationally Sparse Gaussian Processes*, James Hensman, Alexander G. de G. Matthews, Maurizio Filippone, Zoubin Ghahramani

Parameters

- **x** (*torch.Tensor*) – A input data for training. Its first dimension is the number of data points.
- **y** (*torch.Tensor*) – An output data for training. Its last dimension is the number of data points.
- **kernel** (*Kernel*) – A Pyro kernel object, which is the covariance function k .
- **Xu** (*torch.Tensor*) – Initial values for inducing points, which are parameters of our model.
- **Likelihood likelihood** (*likelihood*) – A likelihood object.
- **mean_function** (*callable*) – An optional mean function m of this Gaussian process. By default, we use zero mean.
- **latent_shape** (*torch.Size*) – Shape for latent processes (*batch_shape* of $q(u)$). By default, it equals to output batch shape $y.shape[-1]$. For the multi-class classification problems, `latent_shape[-1]` should correspond to the number of classes.
- **num_data** (*int*) – The size of full training dataset. It is useful for training this model with mini-batch.
- **whiten** (*bool*) – A flag to tell if variational parameters `u_loc` and `u_scale_tril` are transformed by the inverse of `Luu`, where `Luu` is the lower triangular decomposition of $kernel(X_u, X_u)$. Enable this flag will help optimization.
- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.
- **name** (*str*) – Name of this model.

model ()

A “model” stochastic function. If `self.y` is `None`, this method returns mean and variance of the Gaussian Process prior.

guide ()

A “guide” stochastic function to be used in variational inference methods. It also gives posterior information to the method `forward()` for prediction.

forward (*Xnew*, *full_cov=False*)

Computes the mean and covariance matrix (or variance) of Gaussian Process posterior on a test input data X_{new} :

$$p(f^* | X_{new}, X, y, k, X_u, u_{loc}, u_{scale_tril}) = \mathcal{N}(loc, cov).$$

Note: Variational parameters `u_loc`, `u_scale_tril`, the inducing-point parameter X_u , together with kernel's parameters have been learned from a training procedure (MCMC or SVI).

Parameters

- **Xnew** (`torch.Tensor`) – A input data for testing. Note that `Xnew.shape[1:]` must be the same as `self.X.shape[1:]`.
- **full_cov** (`bool`) – A flag to decide if we want to predict full covariance matrix or just variance.

Returns loc and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

10.2.2 Kernels

class Kernel (`input_dim, active_dims=None, name=None`)

Bases: `pyro.contrib.gp.util.Parameterized`

Base class for kernels used in this Gaussian Process module.

Every inherited class should implement a `forward()` pass which takes inputs X , Z and returns their covariance matrix.

To construct a new kernel from the old ones, we can use methods `add()`, `mul()`, `exp()`, `warp()`, `vertical_scale()`.

References:

[1] *Gaussian Processes for Machine Learning*, Carl E. Rasmussen, Christopher K. I. Williams

Parameters

- **input_dim** (`int`) – Number of feature dimensions of inputs.
- **variance** (`torch.Tensor`) – Variance parameter of this kernel.
- **active_dims** (`list`) – List of feature dimensions of the input which the kernel acts on.
- **name** (`str`) – Name of the kernel.

forward (X , $Z=None$, `diag=False`)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- **Z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

add (`other, name=None`)

Creates a new kernel from a sum/direct sum of this kernel object and `other`.

Parameters

- **other** (*Kernel* or *numbers.Number*) – A kernel to be combined with this kernel object.
- **name** (*str*) – An optional name for the derived kernel.

Returns a Sum kernel

Return type *Sum*

mul (*other, name=None*)

Creates a new kernel from a product/tensor product of this kernel object and *other*.

Parameters

- **other** (*Kernel* or *numbers.Number*) – A kernel to be combined with this kernel object.
- **name** (*str*) – An optional name for the derived kernel.

Returns a Product kernel

Return type *Product*

exp (*name=None*)

Creates a new kernel according to

$$k_{new}(x, z) = \exp(k(x, z)).$$

Parameters **name** (*str*) – An optional name for the derived kernel.

Returns an Exponent kernel

Return type *Exponent*

vertical_scale (*vscaling_fn, name=None*)

Creates a new kernel according to

$$k_{new}(x, z) = f(x)k(x, z)f(z),$$

where *f* is a function.

Parameters

- **vscaling_fn** (*callable*) – A vertical scaling function *f*.
- **name** (*str*) – An optional name for the derived kernel.

Returns a vertical scaled kernel

Return type *VerticalScaling*

warp (*iwarping_fn=None, owarping_coef=None, name=None*)

Creates a new kernel according to

$$k_{new}(x, z) = q(k(f(x), f(z))),$$

where *f* is an function and *q* is a polynomial with non-negative coefficients *owarping_coef*.

Parameters

- **iwarping_fn** (*callable*) – An input warping function *f*.
- **owarping_coef** (*list*) – A list of coefficients of the output warping polynomial. These coefficients must be non-negative.
- **name** (*str*) – An optional name for the derived kernel.

Returns a warped kernel

Return type *Warping*

get_subkernel (*name*)

Returns the subkernel corresponding to *name*.

Parameters **name** (*str*) – Name of the subkernel.

Returns A subkernel.

Return type *Kernel*

class Combination (*kern0, kern1, name=None*)

Bases: *pyro.contrib.gp.kernels.kernel.Kernel*

Base class for kernels derived from a combination of kernels.

Parameters

- **kern0** (*Kernel*) – First kernel to combine.
- **kern1** (*Kernel or numbers.Number*) – Second kernel to combine.

class Sum (*kern0, kern1, name=None*)

Bases: *pyro.contrib.gp.kernels.kernel.Combination*

Returns a new kernel which acts like a sum/direct sum of two kernels. The second kernel can be a constant.

forward (*X, Z=None, diag=False*)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times input_dim$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of *X* and *Z* with shape $N \times M$

Return type *torch.Tensor*

class Product (*kern0, kern1, name=None*)

Bases: *pyro.contrib.gp.kernels.kernel.Combination*

Returns a new kernel which acts like a product/tensor product of two kernels. The second kernel can be a constant.

forward (*X, Z=None, diag=False*)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times input_dim$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of *X* and *Z* with shape $N \times M$

Return type *torch.Tensor*

class Transforming (*kern, name=None*)

Bases: *pyro.contrib.gp.kernels.kernel.Kernel*

Base class for kernels derived from a kernel by some transforms such as warping, exponent, vertical scaling.

Parameters `kern` (`Kernel`) – The original kernel.

class `Exponent` (`kern, name=None`)

Bases: `pyro.contrib.gp.kernels.kernel.Transforming`

Creates a new kernel according to

$$k_{new}(x, z) = \exp(k(x, z)).$$

forward (`X, Z=None, diag=False`)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- `x` (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- `z` (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- `diag` (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of `X` and `Z` with shape $N \times M$

Return type `torch.Tensor`

class `VerticalScaling` (`kern, vscaling_fn, name=None`)

Bases: `pyro.contrib.gp.kernels.kernel.Transforming`

Creates a new kernel according to

$$k_{new}(x, z) = f(x)k(x, z)f(z),$$

where `f` is a function.

Parameters `vscaling_fn` (`callable`) – A vertical scaling function `f`.

forward (`X, Z=None, diag=False`)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- `x` (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- `z` (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- `diag` (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of `X` and `Z` with shape $N \times M$

Return type `torch.Tensor`

class `Warping` (`kern, iwarping_fn=None, owarping_coef=None, name=None`)

Bases: `pyro.contrib.gp.kernels.kernel.Transforming`

Creates a new kernel according to

$$k_{new}(x, z) = q(k(f(x), f(z))),$$

where `f` is an function and `q` is a polynomial with non-negative coefficients `owarping_coef`.

We can take advantage of `f` to combine a Gaussian Process kernel with a deep learning architecture. For example:

```
>>> linear = torch.nn.Linear(10, 3)
# register its parameters to Pyro's ParamStore and wrap it by lambda
# to call the primitive pyro.module each time we use the linear function
>>> pyro_linear_fn = lambda x: pyro.module("linear", linear)(x)
>>> kernel = gp.kernels.Matern52(input_dim=3, lengthscale=torch.ones(3))
>>> warped_kernel = gp.kernels.Warping(kernel, pyro_linear_fn)
```

Reference:

[1] *Deep Kernel Learning*, Andrew G. Wilson, Zhiting Hu, Ruslan Salakhutdinov, Eric P. Xing

Parameters

- **iwarping_fn** (*callable*) – An input warping function f .
- **owarping_coef** (*list*) – A list of coefficients of the output warping polynomial. These coefficients must be non-negative.

forward ($X, Z=None, diag=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times input_dim$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class Isotropy (*input_dim, variance=None, lengthscale=None, active_dims=None, name=None*)

Bases: `pyro.contrib.gp.kernels.kernel.Kernel`

Base class for a family of isotropic covariance kernels which are functions of the distance $|x - z|/l$, where l is the length-scale parameter.

By default, the parameter `lengthscale` has size 1. To use the isotropic version (different lengthscale for each dimension), make sure that `lengthscale` has size equal to `input_dim`.

Parameters **lengthscale** (*torch.Tensor*) – Length-scale parameter of this kernel.

class RBF (*input_dim, variance=None, lengthscale=None, active_dims=None, name='RBF'*)

Bases: `pyro.contrib.gp.kernels.isotropic.Isotropy`

Implementation of Radial Basis Function kernel:

$$k(x, z) = \sigma^2 \exp\left(-0.5 \times \frac{|x-z|^2}{l^2}\right).$$

Note: This kernel also has name *Squared Exponential* in literature.

forward ($X, Z=None, diag=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times input_dim$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times input_dim$.

- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class RationalQuadratic (*input_dim, variance=None, lengthscale=None, scale_mixture=None, active_dims=None, name='RationalQuadratic'*)

Bases: `pyro.contrib.gp.kernels.isotropic.Isotropy`

Implementation of RationalQuadratic kernel:

$$k(x, z) = \sigma^2 \left(1 + 0.5 \times \frac{|x-z|^2}{\alpha l^2} \right)^{-\alpha}.$$

Parameters **scale_mixture** (`torch.Tensor`) – Scale mixture (α) parameter of this kernel. Should have size 1.

forward ($X, Z=None, diag=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- **Z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class Exponential (*input_dim, variance=None, lengthscale=None, active_dims=None, name='Exponential'*)

Bases: `pyro.contrib.gp.kernels.isotropic.Isotropy`

Implementation of Exponential kernel:

$$k(x, z) = \sigma^2 \exp\left(-\frac{|x-z|}{l}\right).$$

forward ($X, Z=None, diag=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- **Z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class Matern32 (*input_dim, variance=None, lengthscale=None, active_dims=None, name='Matern32'*)

Bases: `pyro.contrib.gp.kernels.isotropic.Isotropy`

Implementation of Matern32 kernel:

$$k(x, z) = \sigma^2 \left(1 + \sqrt{3} \times \frac{|x-z|}{l} \right) \exp\left(-\sqrt{3} \times \frac{|x-z|}{l}\right).$$

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type *torch.Tensor*

class Matern52 (*input_dim, variance=None, lengthscale=None, active_dims=None, name='Matern52'*)

Bases: *pyro.contrib.gp.kernels.isotropic.Isotropy*

Implementation of Matern52 kernel:

$$k(x, z) = \sigma^2 \left(1 + \sqrt{5} \times \frac{|x-z|}{l} + \frac{5}{3} \times \frac{|x-z|^2}{l^2} \right) \exp \left(-\sqrt{5} \times \frac{|x-z|}{l} \right).$$

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type *torch.Tensor*

class Cosine (*input_dim, variance=None, lengthscale=None, active_dims=None, name='Cosine'*)

Bases: *pyro.contrib.gp.kernels.isotropic.Isotropy*

Implementation of Cosine kernel:

$$k(x, z) = \sigma^2 \cos \left(\frac{|x-z|}{l} \right).$$

Parameters **lengthscale** (*torch.Tensor*) – Length-scale parameter of this kernel.

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type *torch.Tensor*

class Periodic (*input_dim*, *variance=None*, *lengthscale=None*, *period=None*, *active_dims=None*, *name='Periodic'*)

Bases: `pyro.contrib.gp.kernels.kernel.Kernel`

Implementation of Periodic kernel:

$$k(x, z) = \sigma^2 \exp\left(-2 \times \frac{\sin^2(\pi(x-z)/p)}{l^2}\right),$$

where p is the period parameter.

References:

[1] *Introduction to Gaussian processes*, David J.C. MacKay

Parameters

- **lengthscale** (`torch.Tensor`) – Length scale parameter of this kernel.
- **period** (`torch.Tensor`) – Period parameter of this kernel.

forward (*X*, *Z=None*, *diag=False*)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (`torch.Tensor`) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class Constant (*input_dim*, *variance=None*, *active_dims=None*, *name='Constant'*)

Bases: `pyro.contrib.gp.kernels.kernel.Kernel`

Implementation of Constant kernel:

$$k(x, z) = \sigma^2.$$

forward (*X*, *Z=None*, *diag=False*)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (`torch.Tensor`) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class WhiteNoise (*input_dim*, *variance=None*, *active_dims=None*, *name='WhiteNoise'*)

Bases: `pyro.contrib.gp.kernels.kernel.Kernel`

Implementation of WhiteNoise kernel:

$$k(x, z) = \sigma^2 \delta(x, z),$$

where δ is a Dirac delta function.

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type *torch.Tensor*

class **DotProduct** (*input_dim, variance=None, active_dims=None, name=None*)

Bases: *pyro.contrib.gp.kernels.kernel.Kernel*

Base class for kernels which are functions of $x \cdot z$.

class **Linear** (*input_dim, variance=None, active_dims=None, name='Linear'*)

Bases: *pyro.contrib.gp.kernels.dot_product.DotProduct*

Implementation of Linear kernel:

$$k(x, z) = \sigma^2 x \cdot z.$$

Doing Gaussian Process regression with linear kernel is equivalent to doing a linear regression.

Note: Here we implement the homogeneous version. To use the inhomogeneous version, consider using *Polynomial* kernel with `degree=1` or making a *Sum* with a *Bias* kernel.

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **X** (*torch.Tensor*) – A 2D tensor with shape $N \times \text{input_dim}$.
- **Z** (*torch.Tensor*) – An (optional) 2D tensor with shape $M \times \text{input_dim}$.
- **diag** (*bool*) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type *torch.Tensor*

class **Polynomial** (*input_dim, variance=None, bias=None, degree=1, active_dims=None, name='Polynomial'*)

Bases: *pyro.contrib.gp.kernels.dot_product.DotProduct*

Implementation of Polynomial kernel:

$$k(x, z) = \sigma^2 (\text{bias} + x \cdot z)^d.$$

Parameters

- **bias** (*torch.Tensor*) – Bias parameter of this kernel. Should be positive.
- **degree** (*int*) – Degree d of the polynomial.

forward ($X, Z=None, \text{diag}=False$)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **x** (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- **z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

class `Brownian` (`input_dim`, `variance=None`, `active_dims=None`, `name='Brownian'`)

Bases: `pyro.contrib.gp.kernels.kernel.Kernel`

This kernel corresponds to a two-sided Brownian motion (Wiener process):

$$k(x, z) = \begin{cases} \sigma^2 \min(|x|, |z|), & \text{if } x \cdot z \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

Note that the input dimension of this kernel must be 1.

Reference:

[1] *Theory and Statistical Applications of Stochastic Processes*, Yuliya Mishura, Georgiy Shevchenko

forward (X , $Z=None$, `diag=False`)

Calculates covariance matrix of inputs on active dimensionals.

Parameters

- **x** (`torch.Tensor`) – A 2D tensor with shape $N \times input_dim$.
- **z** (`torch.Tensor`) – An (optional) 2D tensor with shape $M \times input_dim$.
- **diag** (`bool`) – A flag to decide if we want to return full covariance matrix or just its diagonal part.

Returns covariance matrix of X and Z with shape $N \times M$

Return type `torch.Tensor`

10.2.3 Likelihoods

class `Likelihood` (`name=None`)

Bases: `pyro.contrib.gp.util.Parameterized`

Base class for likelihoods used in Gaussian Process.

Every inherited class should implement a forward pass which takes an input f and returns a sample y .

forward (f_loc , f_var , $y=None$)

Samples y given f_loc , f_var .

Parameters

- **f_loc** (`torch.Tensor`) – Mean of latent function output.
- **f_var** (`torch.Tensor`) – Variance of latent function output.
- **y** (`torch.Tensor`) – Training output tensor.

Returns a tensor sampled from likelihood

Return type `torch.Tensor`

```
class Gaussian (variance=None, name='Gaussian')
    Bases: pyro.contrib.gp.likelihoods.likelihood.Likelihood
```

Implementation of Gaussian likelihood, which is used for regression problems.

Gaussian likelihood uses Normal distribution.

Parameters **variance** (`torch.Tensor`) – A variance parameter, which plays the role of noise in regression problems.

```
forward (f_loc, f_var, y=None)
    Samples  $y$  given  $f_{loc}$ ,  $f_{var}$  according to
```

$$y \sim \mathcal{N}(f_{loc}, f_{var} + \epsilon),$$

where ϵ is the variance parameter of this likelihood.

Parameters

- **f_loc** (`torch.Tensor`) – Mean of latent function output.
- **f_var** (`torch.Tensor`) – Variance of latent function output.
- **y** (`torch.Tensor`) – Training output tensor.

Returns a tensor sampled from likelihood

Return type `torch.Tensor`

```
class Binary (response_function=None, name='Binary')
    Bases: pyro.contrib.gp.likelihoods.likelihood.Likelihood
```

Implementation of Binary likelihood, which is used for binary classification problems.

Binary likelihood uses Bernoulli distribution, so the output of `response_function` should be in range (0, 1). By default, we use *sigmoid* function.

Parameters **response_function** (`callable`) – A mapping to correct domain for Binary likelihood.

```
forward (f_loc, f_var, y=None)
    Samples  $y$  given  $f_{loc}$ ,  $f_{var}$  according to
```

$$f \sim \mathcal{N}(f_{loc}, f_{var}),$$
$$y \sim \mathbb{B}(f).$$

Note: The log likelihood is estimated using Monte Carlo with 1 sample of f .

Parameters

- **f_loc** (`torch.Tensor`) – Mean of latent function output.
- **f_var** (`torch.Tensor`) – Variance of latent function output.
- **y** (`torch.Tensor`) – Training output tensor.

Returns a tensor sampled from likelihood

Return type `torch.Tensor`

class MultiClass (*num_classes*, *response_function=None*, *name='MultiClass'*)
 Bases: `pyro.contrib.gp.likelihoods.likelihood.Likelihood`

Implementation of MultiClass likelihood, which is used for multi-class classification problems.

MultiClass likelihood uses `Categorical` distribution, so `response_function` should normalize its input's rightmost axis. By default, we use `softmax` function.

Parameters

- **num_classes** (*int*) – Number of classes for prediction.
- **response_function** (*callable*) – A mapping to correct domain for MultiClass likelihood.

forward (*f_loc*, *f_var*, *y=None*)
 Samples *y* given *f_loc*, *f_var* according to

$$f \sim \mathbb{N}(f_{loc}, f_{var}),$$

$$y \sim \text{Categorical}(f).$$

Note: The log likelihood is estimated using Monte Carlo with 1 sample of *f*.

Parameters

- **f_loc** (*torch.Tensor*) – Mean of latent function output.
- **f_var** (*torch.Tensor*) – Variance of latent function output.
- **y** (*torch.Tensor*) – Training output tensor.

Returns a tensor sampled from likelihood

Return type `torch.Tensor`

class Poisson (*response_function=None*, *name='Poisson'*)
 Bases: `pyro.contrib.gp.likelihoods.likelihood.Likelihood`

Implementation of Poisson likelihood, which is used for count data.

Poisson likelihood uses the `Poisson` distribution, so the output of `response_function` should be positive. By default, we use `torch.exp()` as response function, corresponding to a log-Gaussian Cox process.

Parameters **response_function** (*callable*) – A mapping to positive real numbers.

forward (*f_loc*, *f_var*, *y=None*)
 Samples *y* given *f_loc*, *f_var* according to

$$f \sim \mathbb{N}(f_{loc}, f_{var}),$$

$$y \sim \text{Poisson}(\exp(f)).$$

Note: The log likelihood is estimated using Monte Carlo with 1 sample of *f*.

Parameters

- **f_loc** (*torch.Tensor*) – Mean of latent function output.

- **f_var** (*torch.Tensor*) – Variance of latent function output.
- **y** (*torch.Tensor*) – Training output tensor.

Returns a tensor sampled from likelihood

Return type `torch.Tensor`

Util

class Parameterized (*name=None*)

Bases: `torch.nn.modules.module.Module`

Base class for other modules in Gaussian Process module.

Parameters of this object can be set priors, set constraints, or fixed to a specific value.

By default, data of a parameter is a float `torch.Tensor` (unless we use `torch.set_default_tensor_type()` to change default tensor type). To cast these parameters to a correct data type or GPU device, we can call methods such as `double()` or `cuda()`. See `torch.nn.Module` for more information.

Parameters **name** (*str*) – Name of this object.

set_prior (*param, prior*)

Sets a prior to a parameter.

Parameters

- **param** (*str*) – Name of the parameter.
- **prior** (*Distribution*) – A Pyro prior distribution.

set_constraint (*param, constraint*)

Sets a constraint to a parameter.

Parameters

- **param** (*str*) – Name of the parameter.
- **constraint** (*Constraint*) – A PyTorch constraint. See `torch.distributions.constraints` for a list of constraints.

fix_param (*param, value=None*)

Fixes a parameter to a specific value. If `value=None`, fixes the parameter to the default value.

Parameters

- **param** (*str*) – Name of the parameter.
- **value** (*torch.Tensor*) – Fixed value.

set_mode (*mode*)

Sets mode of this object to be able to use its parameters in stochastic functions. If `mode="model"`, a parameter with prior will get its value from the primitive `pyro.sample()`. If `mode="guide"` or there is no prior on a parameter, `pyro.param()` will be called.

This method automatically sets mode for submodules which belong to `Parameterized` class.

Parameters **mode** (*str*) – Either “model” or “guide”.

get_param (*param*)

Gets the current value of a parameter. The correct behavior will depend on mode of this object (see `set_mode()` method).

Parameters `param` (*str*) – Name of the parameter.

conditional (*Xnew*, *X*, *kernel*, *f_loc*, *f_scale_tril=None*, *Lff=None*, *full_cov=False*, *whiten=False*, *jitter=1e-06*)

Given X_{new} , predicts loc and covariance matrix of the conditional multivariate normal distribution

$$p(f^*(X_{new}) \mid X, k, f_{loc}, f_{scale_tril}).$$

Here `f_loc` and `f_scale_tril` are variation parameters of the variational distribution

$$q(f \mid f_{loc}, f_{scale_tril}) \sim p(f \mid X, y),$$

where f is the function value of the Gaussian Process given input X

$$p(f(X)) \sim \mathcal{N}(0, k(X, X))$$

and y is computed from f by some likelihood function $p(y \mid f)$.

In case `f_scale_tril=None`, we consider $f = f_{loc}$ and computes

$$p(f^*(X_{new}) \mid X, k, f).$$

In case `f_scale_tril` is not `None`, we follow the derivation from reference [1]. For the case `f_scale_tril=None`, we follow the popular reference [2].

References:

[1] [Sparse GPs: approximate the posterior, not the model](#)

[2] *Gaussian Processes for Machine Learning*, Carl E. Rasmussen, Christopher K. I. Williams

Parameters

- **Xnew** (*torch.Tensor*) – A new input data.
- **X** (*torch.Tensor*) – An input data to be conditioned on.
- **kernel** (*Kernel*) – A Pyro kernel object.
- **f_loc** (*torch.Tensor*) – Mean of $q(f)$. In case `f_scale_tril=None`, $f_{loc} = f$.
- **f_scale_tril** (*torch.Tensor*) – Lower triangular decomposition of covariance matrix of $q(f)$'s.
- **Lff** (*torch.Tensor*) – Lower triangular decomposition of $kernel(X, X)$ (optional).
- **full_cov** (*bool*) – A flag to decide if we want to return full covariance matrix or just variance.
- **whiten** (*bool*) – A flag to tell if `f_loc` and `f_scale_tril` are already transformed by the inverse of `Lff`.
- **jitter** (*float*) – A small positive term which is added into the diagonal part of a covariance matrix to help stabilize its Cholesky decomposition.

Returns loc and covariance matrix (or variance) of $p(f^*(X_{new}))$

Return type `tuple(torch.Tensor, torch.Tensor)`

10.3 Named Data Structures

The `pyro.contrib.named` module is a thin syntactic layer on top of Pyro. It allows Pyro models to be written to look like programs with operating on Python data structures like `latent.x.sample_(...)`, rather than programs with string-labeled statements like `x = pyro.sample("x", ...)`.

This module provides three container data structures named `Object`, `List`, and `Dict`. These data structures are intended to be nested in each other. Together they track the address of each piece of data in each data structure, so that this address can be used as a Pyro site. For example:

```
>>> state = named.Object("state")
>>> print(str(state))
state

>>> z = state.x.y.z # z is just a placeholder.
>>> print(str(z))
state.x.y.z

>>> state.xs = named.List() # Create a contained list.
>>> x0 = state.xs.add()
>>> print(str(x0))
state.xs[0]

>>> state.ys = named.Dict()
>>> foo = state.ys['foo']
>>> print(str(foo))
state.ys['foo']
```

These addresses can now be used inside `sample`, `observe` and `param` statements. These named data structures even provide in-place methods that alias Pyro statements. For example:

```
>>> state = named.Object("state")
>>> loc = state.loc.param_(torch.zeros(1, requires_grad=True))
>>> scale = state.scale.param_(torch.ones(1, requires_grad=True))
>>> z = state.z.sample_(dist.normal, loc, scale)
>>> state.x.observe_(dist.normal, z, loc, scale)
```

For deeper examples of how these can be used in model code, see the [Tree Data](#) and [Mixture](#) examples.

Authors: Fritz Obermeyer, Alexander Rush

class `Object` (*name*)

Bases: `object`

Object to hold immutable latent state.

This object can serve either as a container for nested latent state or as a placeholder to be replaced by a tensor via a `named.sample`, `named.observe`, or `named.param` statement. When used as a placeholder, `Object` objects take the place of strings in normal `pyro.sample` statements.

Parameters `name` (*str*) – The name of the object.

Example:

```
state = named.Object("state")
state.x = 0
state.ys = named.List()
state.zs = named.Dict()
state.a.b.c.d.e.f.g = 0 # Creates a chain of named.Objects.
```

Warning: This data structure is write-once: data may be added but may not be mutated or removed. Trying to mutate this data structure may result in silent errors.

sample_(*fn*, **args*, ***kwargs*)

Calls the stochastic function *fn* with additional side-effects depending on *name* and the enclosing context (e.g. an inference algorithm). See [Intro I](#) and [Intro II](#) for a discussion.

Parameters

- **name** – name of sample
- **fn** – distribution class or function
- **obs** – observed datum (optional; should only be used in context of inference) optionally specified in *kwargs*
- **infer** (*dict*) – Optional dictionary of inference parameters specified in *kwargs*. See inference documentation for details.

Returns sample

param_(**args*, ***kwargs*)

Saves the variable as a parameter in the param store. To interact with the param store or write to disk, see [Parameters](#).

Parameters **name** – name of parameter

Returns parameter

class List (*name=None*)

Bases: `list`

List-like object to hold immutable latent state.

This must either be given a name when constructed:

```
latent = named.List("root")
```

or must be immediately stored in a `named.Object`:

```
latent = named.Object("root")
latent.xs = named.List() # Must be bound to a Object before use.
```

Warning: This data structure is write-once: data may be added but may not be mutated or removed. Trying to mutate this data structure may result in silent errors.

add()

Append one new `named.Object`.

Returns a new latent object at the end

Return type *named.Object*

class Dict (*name=None*)

Bases: `dict`

Dict-like object to hold immutable latent state.

This must either be given a name when constructed:

```
latent = named.Dict("root")
```

or must be immediately stored in a `named.Object`:

```
latent = named.Object("root")
latent.xs = named.Dict() # Must be bound to a Object before use.
```

Warning: This data structure is write-once: data may be added but may not be mutated or removed. Trying to mutate this data structure may result in silent errors.

CHAPTER 11

Indices and tables

- `genindex`
- `search`

p

pyro.contrib.autoguide, 51
pyro.contrib.gp, 54
pyro.contrib.gp.kernels.brownian, 73
pyro.contrib.gp.kernels.dot_product, 72
pyro.contrib.gp.kernels.isotropic, 68
pyro.contrib.gp.kernels.kernel, 64
pyro.contrib.gp.kernels.periodic, 70
pyro.contrib.gp.kernels.static, 71
pyro.contrib.gp.likelihoods.binary, 74
pyro.contrib.gp.likelihoods.gaussian,
73
pyro.contrib.gp.likelihoods.likelihood,
73
pyro.contrib.gp.likelihoods.multi_class,
75
pyro.contrib.gp.likelihoods.poisson, 75
pyro.contrib.gp.models.gpr, 57
pyro.contrib.gp.models.model, 54
pyro.contrib.gp.models.sgpr, 59
pyro.contrib.gp.models.vgp, 61
pyro.contrib.gp.models.vsgp, 62
pyro.contrib.gp.util, 76
pyro.contrib.named, 77
pyro.distributions.torch, 17
pyro.infer.abstract_infer, 12
pyro.infer.elbo, 10
pyro.infer.importance, 12
pyro.infer.svi, 9
pyro.infer.trace_elbo, 10
pyro.infer.traceenum_elbo, 11
pyro.infer.tracegraph_elbo, 11
pyro.nn.auto_reg_nn, 33
pyro.optim.optim, 35
pyro.optim.pytorch_optimizers, 36
pyro.params.param_store, 31
pyro.poutine.block_messenger, 45
pyro.poutine.condition_messenger, 46
pyro.poutine.escape_messenger, 46
pyro.poutine.handlers, 37
pyro.poutine.indep_messenger, 46
pyro.poutine.lift_messenger, 46
pyro.poutine.messenger, 45
pyro.poutine.replay_messenger, 47
pyro.poutine.runtime, 48
pyro.poutine.scale_messenger, 47
pyro.poutine.trace_messenger, 47
pyro.poutine.util, 48

Symbols

`__call__()` (AutoContinuous method), 52
`__call__()` (AutoDiscreteParallel method), 54
`__call__()` (AutoGuide method), 51
`__call__()` (AutoGuideList method), 52
`__call__()` (Distribution method), 21

A

Adadelta() (in module `pyro.optim.pytorch_optimizers`), 36
 Adagrad() (in module `pyro.optim.pytorch_optimizers`), 36
 AdagradRMSProp() (in module `pyro.optim.optim`), 35
 Adam() (in module `pyro.optim.pytorch_optimizers`), 36
 Adamax() (in module `pyro.optim.pytorch_optimizers`), 36
 add() (AutoGuideList method), 52
 add() (Empirical method), 24
 add() (Kernel method), 64
 add() (List method), 79
 add_edge (Trace attribute), 43
 add_node() (Trace method), 43
 all_escape() (in module `pyro.poutine.util`), 48
 am_i_wrapped() (in module `pyro.poutine.runtime`), 48
 apply_stack() (in module `pyro.poutine.runtime`), 48
 arg_constraints (Binomial attribute), 23
 arg_constraints (Delta attribute), 24
 arg_constraints (Empirical attribute), 25
 arg_constraints (HalfCauchy attribute), 25
 arg_constraints (LowRankMultivariateNormal attribute), 26
 arg_constraints (OMTMultivariateNormal attribute), 27
 arg_constraints (VonMises attribute), 28
 arn (InverseAutoregressiveFlow attribute), 29
 ASGD() (in module `pyro.optim.pytorch_optimizers`), 36
 AutoContinuous (class in `pyro.contrib.autoguide`), 52
 AutoDiagonalNormal (class in `pyro.contrib.autoguide`), 53
 AutoDiscreteParallel (class in `pyro.contrib.autoguide`), 54
 AutoGuide (class in `pyro.contrib.autoguide`), 51

AutoGuideList (class in `pyro.contrib.autoguide`), 52
 AutoMultivariateNormal (class in `pyro.contrib.autoguide`), 53
 AutoRegressiveNN (class in `pyro.nn.auto_reg_nn`), 33

B

Bernoulli (class in `pyro.distributions`), 17
 Beta (class in `pyro.distributions`), 17
 Binary (class in `pyro.contrib.gp.likelihoods.binary`), 74
 Binomial (class in `pyro.distributions`), 23
 block() (in module `pyro.poutine`), 38
 BlockMessenger (class in `pyro.poutine.block_messenger`), 45
 Brownian (class in `pyro.contrib.gp.kernels.brownian`), 73

C

Categorical (class in `pyro.distributions`), 17
 Cauchy (class in `pyro.distributions`), 17
 Chi2 (class in `pyro.distributions`), 18
 cleanup() (HMC method), 14
 clear() (ParamStoreDict method), 31
 clear_param_store() (in module `pyro`), 8
 ClippedAdam() (in module `pyro.optim.optim`), 35
 Combination (class in `pyro.contrib.gp.kernels.kernel`), 66
 compute_log_prob() (Trace method), 43
 compute_score_parts() (Trace method), 43
 CondIndepStackFrame (class in `pyro.poutine.indep_messenger`), 46
 condition() (in module `pyro.poutine`), 38
 conditional() (in module `pyro.contrib.gp.util`), 77
 ConditionMessenger (class in `pyro.poutine.condition_messenger`), 46
 Constant (class in `pyro.contrib.gp.kernels.static`), 71
 copy() (Trace method), 43
 Cosine (class in `pyro.contrib.gp.kernels.periodic`), 70

D

default_process_message() (in module `pyro.poutine.runtime`), 48

Delta (class in pyro.distributions), 24
 diagnostics() (HMC method), 14
 Dict (class in pyro.contrib.named), 79
 Dirichlet (class in pyro.distributions), 18
 discrete_escape() (in module pyro.poutine.util), 48
 Distribution (class in pyro.distributions), 20
 do() (in module pyro.poutine), 39
 DotProduct (class in pyro.contrib.gp.kernels.dot_product),
 72

E

edges (Trace attribute), 43
 ELBO (class in pyro.infer.elbo), 10
 Empirical (class in pyro.distributions), 24
 EmpiricalMarginal (class in pyro.infer.abstract_infer), 12
 enable_validation() (in module pyro), 8
 end_warmup() (HMC method), 14
 entropy() (HalfCauchy method), 25
 enum() (in module pyro.poutine), 39
 enum_extend() (in module pyro.poutine.util), 49
 enumerate_support() (Binomial method), 23
 enumerate_support() (Distribution method), 21
 enumerate_support() (Empirical method), 25
 escape() (in module pyro.poutine), 39
 EscapeMessenger (class in pyro.poutine.escape_messenger), 46
 evaluate_loss() (SVI method), 9
 event_shape (Empirical attribute), 25
 exp() (Kernel method), 65
 Exponent (class in pyro.contrib.gp.kernels.kernel), 67
 Exponential (class in pyro.contrib.gp.kernels.isotropic),
 69
 Exponential (class in pyro.distributions), 18
 ExponentialFamily (class in pyro.distributions), 18

F

FisherSnedecor (class in pyro.distributions), 18
 fix_param() (Parameterized method), 76
 forward() (AutoRegressiveNN method), 33
 forward() (Binary method), 74
 forward() (Brownian method), 73
 forward() (Constant method), 71
 forward() (Cosine method), 70
 forward() (Exponent method), 67
 forward() (Exponential method), 69
 forward() (Gaussian method), 74
 forward() (GPModel method), 56
 forward() (GPRegression method), 58
 forward() (Kernel method), 64
 forward() (Likelihood method), 73
 forward() (Linear method), 72
 forward() (MaskedLinear method), 34
 forward() (Matern32 method), 69
 forward() (Matern52 method), 70

forward() (MultiClass method), 75
 forward() (Periodic method), 71
 forward() (Poisson method), 75
 forward() (Polynomial method), 72
 forward() (Product method), 66
 forward() (RationalQuadratic method), 69
 forward() (RBF method), 68
 forward() (SparseGPRegression method), 60
 forward() (Sum method), 66
 forward() (VariationalGP method), 62
 forward() (VariationalSparseGP method), 63
 forward() (VerticalScaling method), 67
 forward() (Warping method), 68
 forward() (WhiteNoise method), 71

G

Gamma (class in pyro.distributions), 18
 Gaussian (class in pyro.contrib.gp.likelihoods.gaussian),
 73
 Geometric (class in pyro.distributions), 18
 get_all_param_names() (ParamStoreDict method), 31
 get_mask_encoding() (AutoRegressiveNN method), 33
 get_param() (Parameterized method), 76
 get_param() (ParamStoreDict method), 32
 get_param_store() (in module pyro), 7
 get_permutation() (AutoRegressiveNN method), 34
 get_samples_and_weights() (Empirical method), 25
 get_state() (ParamStoreDict method), 32
 get_state() (PyroOptim method), 35
 get_subkernel() (Kernel method), 66
 get_trace() (TraceHandler method), 47
 get_trace() (TraceMessenger method), 47
 GPModel (class in pyro.contrib.gp.models.model), 54
 GPRegression (class in pyro.contrib.gp.models.gpr), 57
 graph (Trace attribute), 44
 guide() (GPModel method), 56
 guide() (GPRegression method), 58
 guide() (SparseGPRegression method), 60
 guide() (VariationalGP method), 62
 guide() (VariationalSparseGP method), 63
 Gumbel (class in pyro.distributions), 18

H

HalfCauchy (class in pyro.distributions), 25
 Handler (class in pyro.poutine.messenger), 45
 has_enumerate_support (Binomial attribute), 23
 has_enumerate_support (Distribution attribute), 21
 has_enumerate_support (Empirical attribute), 25
 has_rsamle (Delta attribute), 24
 has_rsamle (Distribution attribute), 21
 has_rsamle (LowRankMultivariateNormal attribute), 26
 has_rsamle (Rejector attribute), 27
 HMC (class in pyro.infer.mcmc), 13

I

iarange (class in pyro), 6
 identify_dense_edges() (in module pyro.poutine.trace_messenger), 47
 Importance (class in pyro.infer.importance), 12
 in_degree (Trace attribute), 44
 indep() (in module pyro.poutine), 40
 Independent (class in pyro.distributions), 18
 IndepMessenger (class in pyro.poutine.indep_messenger), 46
 infer_config() (in module pyro.poutine), 40
 initial_trace() (HMC method), 14
 InverseAutoregressiveFlow (class in pyro.distributions), 28
 irange (class in pyro), 6
 is_directed (Trace attribute), 44
 Isotropy (class in pyro.contrib.gp.kernels.isotropic), 68
 iter_stochastic_nodes() (Trace method), 44

K

Kernel (class in pyro.contrib.gp.kernels.kernel), 64

L

Laplace (class in pyro.distributions), 19
 LBFGS() (in module pyro.optim.pytorch_optimizers), 36
 lift() (in module pyro.poutine), 40
 LiftMessenger (class in pyro.poutine.lift_messenger), 46
 Likelihood (class in pyro.contrib.gp.likelihoods.likelihood), 73
 Linear (class in pyro.contrib.gp.kernels.dot_product), 72
 List (class in pyro.contrib.named), 79
 load() (ParamStoreDict method), 32
 load() (PyroOptim method), 35
 loc (HalfCauchy attribute), 25
 log_abs_det_jacobian() (InverseAutoregressiveFlow method), 29
 log_prob() (Binomial method), 23
 log_prob() (Delta method), 24
 log_prob() (Distribution method), 21
 log_prob() (Empirical method), 25
 log_prob() (HalfCauchy method), 25
 log_prob() (LowRankMultivariateNormal method), 26
 log_prob() (Rejector method), 27
 log_prob() (VonMises method), 28
 log_prob_sum() (Trace method), 44
 LogisticNormal (class in pyro.distributions), 19
 logits (Binomial attribute), 23
 LogNormal (class in pyro.distributions), 19
 loss() (Trace_ELBO method), 10
 loss() (TraceEnum_ELBO method), 11
 loss() (TraceGraph_ELBO method), 11
 loss_and_grads() (Trace_ELBO method), 10
 loss_and_grads() (TraceEnum_ELBO method), 11

loss_and_grads() (TraceGraph_ELBO method), 11
 LowRankMultivariateNormal (class in pyro.distributions), 26

M

MaskedLinear (class in pyro.nn.auto_reg_nn), 34
 Matern32 (class in pyro.contrib.gp.kernels.isotropic), 69
 Matern52 (class in pyro.contrib.gp.kernels.isotropic), 70
 mc_extend() (in module pyro.poutine.util), 49
 MCMC (class in pyro.infer.mcmc), 13
 mean (Binomial attribute), 23
 mean (Delta attribute), 24
 mean (Empirical attribute), 25
 mean (LowRankMultivariateNormal attribute), 26
 median() (AutoDiagonalNormal method), 54
 median() (AutoMultivariateNormal method), 53
 Messenger (class in pyro.poutine.messenger), 45
 model() (GPModel method), 56
 model() (GPRegression method), 58
 model() (SparseGPRegression method), 60
 model() (VariationalGP method), 62
 model() (VariationalSparseGP method), 63
 module() (in module pyro), 5
 module_from_param_with_module_name() (in module pyro.params.param_store), 32
 mul() (Kernel method), 65
 MultiClass (class in pyro.contrib.gp.likelihoods.multi_class), 75
 Multinomial (class in pyro.distributions), 19
 MultivariateNormal (class in pyro.distributions), 19

N

named_parameters() (ParamStoreDict method), 32
 next_context() (IndepMessenger method), 46
 nodes (Trace attribute), 44
 NonlocalExit, 48
 nonreparam_stochastic_nodes (Trace attribute), 44
 Normal (class in pyro.distributions), 19
 NUTS (class in pyro.infer.mcmc), 15

O

Object (class in pyro.contrib.named), 78
 observation_nodes (Trace attribute), 44
 OMTMultivariateNormal (class in pyro.distributions), 27
 OneHotCategorical (class in pyro.distributions), 19
 optimize() (GPModel method), 57

P

param() (in module pyro), 5
 param_() (Object method), 79
 param_name() (ParamStoreDict method), 32
 param_nodes (Trace attribute), 44
 param_shape (Binomial attribute), 23

[param_with_module_name\(\)](#) (in module `pyro.params.param_store`), 32
[Parameterized](#) (class in `pyro.contrib.gp.util`), 76
[ParamStoreDict](#) (class in `pyro.params.param_store`), 31
[Pareto](#) (class in `pyro.distributions`), 19
[Periodic](#) (class in `pyro.contrib.gp.kernels.periodic`), 70
[Poisson](#) (class in `pyro.contrib.gp.likelihoods.poisson`), 75
[Poisson](#) (class in `pyro.distributions`), 20
[Polynomial](#) (class in `pyro.contrib.gp.kernels.dot_product`), 72
[probs](#) (Binomial attribute), 23
[Product](#) (class in `pyro.contrib.gp.kernels.kernel`), 66
[prune_subsample_sites\(\)](#) (in module `pyro.poutine.util`), 49
[pyro.contrib.autoguide](#) (module), 51
[pyro.contrib.gp](#) (module), 54
[pyro.contrib.gp.kernels.brownian](#) (module), 73
[pyro.contrib.gp.kernels.dot_product](#) (module), 72
[pyro.contrib.gp.kernels.isotropic](#) (module), 68
[pyro.contrib.gp.kernels.kernel](#) (module), 64
[pyro.contrib.gp.kernels.periodic](#) (module), 70
[pyro.contrib.gp.kernels.static](#) (module), 71
[pyro.contrib.gp.likelihoods.binary](#) (module), 74
[pyro.contrib.gp.likelihoods.gaussian](#) (module), 73
[pyro.contrib.gp.likelihoods.likelihood](#) (module), 73
[pyro.contrib.gp.likelihoods.multi_class](#) (module), 75
[pyro.contrib.gp.likelihoods.poisson](#) (module), 75
[pyro.contrib.gp.models.gpr](#) (module), 57
[pyro.contrib.gp.models.model](#) (module), 54
[pyro.contrib.gp.models.sgpr](#) (module), 59
[pyro.contrib.gp.models.vgp](#) (module), 61
[pyro.contrib.gp.models.vsgp](#) (module), 62
[pyro.contrib.gp.util](#) (module), 76
[pyro.contrib.named](#) (module), 77
[pyro.distributions.torch](#) (module), 17
[pyro.infer.abstract_infer](#) (module), 12
[pyro.infer.elbo](#) (module), 10
[pyro.infer.importance](#) (module), 12
[pyro.infer.svi](#) (module), 9
[pyro.infer.trace_elbo](#) (module), 10
[pyro.infer.traceenum_elbo](#) (module), 11
[pyro.infer.tracegraph_elbo](#) (module), 11
[pyro.nn.auto_reg_nn](#) (module), 33
[pyro.optim.optim](#) (module), 35
[pyro.optim.pytorch_optimizers](#) (module), 36
[pyro.params.param_store](#) (module), 31
[pyro.poutine.block_messenger](#) (module), 45
[pyro.poutine.condition_messenger](#) (module), 46
[pyro.poutine.escape_messenger](#) (module), 46
[pyro.poutine.handlers](#) (module), 37
[pyro.poutine.indep_messenger](#) (module), 46
[pyro.poutine.lift_messenger](#) (module), 46
[pyro.poutine.messenger](#) (module), 45
[pyro.poutine.replay_messenger](#) (module), 47

[pyro.poutine.runtime](#) (module), 48
[pyro.poutine.scale_messenger](#) (module), 47
[pyro.poutine.trace_messenger](#) (module), 47
[pyro.poutine.util](#) (module), 48
[PyroOptim](#) (class in `pyro.optim.optim`), 35

Q

[quantiles\(\)](#) (`AutoDiagonalNormal` method), 54
[quantiles\(\)](#) (`AutoMultivariateNormal` method), 53
[queue\(\)](#) (in module `pyro.poutine`), 41

R

[random_module\(\)](#) (in module `pyro`), 5
[RationalQuadratic](#) (class in `pyro.contrib.gp.kernels.isotropic`), 69
[RBF](#) (class in `pyro.contrib.gp.kernels.isotropic`), 68
[Rejector](#) (class in `pyro.distributions`), 27
[RelaxedBernoulli](#) (class in `pyro.distributions`), 20
[RelaxedOneHotCategorical](#) (class in `pyro.distributions`), 20
[remove_node](#) (`Trace` attribute), 44
[reparameterized_nodes](#) (`Trace` attribute), 44
[replace_param\(\)](#) (`ParamStoreDict` method), 32
[replay\(\)](#) (in module `pyro.poutine`), 40
[ReplayMessenger](#) (class in `pyro.poutine.replay_messenger`), 47
[reset_stack\(\)](#) (`NonlocalExit` method), 48
[RMSprop\(\)](#) (in module `pyro.optim.pytorch_optimizers`), 36
[Rprop\(\)](#) (in module `pyro.optim.pytorch_optimizers`), 36
[rsample\(\)](#) (`Delta` method), 24
[rsample\(\)](#) (`LowRankMultivariateNormal` method), 26
[rsample\(\)](#) (`OMTMultivariateNormal` method), 27
[rsample\(\)](#) (`Rejector` method), 27
[run\(\)](#) (`TracePosterior` method), 12

S

[sample\(\)](#) (Binomial method), 23
[sample\(\)](#) (Distribution method), 21
[sample\(\)](#) (Empirical method), 25
[sample\(\)](#) (HMC method), 14
[sample\(\)](#) (in module `pyro`), 5
[sample\(\)](#) (NUTS method), 16
[sample_\(\)](#) (Object method), 79
[sample_latent\(\)](#) (`AutoContinuous` method), 52
[sample_latent\(\)](#) (`AutoDiagonalNormal` method), 54
[sample_latent\(\)](#) (`AutoGuide` method), 52
[sample_latent\(\)](#) (`AutoMultivariateNormal` method), 53
[sample_size](#) (Empirical attribute), 25
[save\(\)](#) (`ParamStoreDict` method), 32
[save\(\)](#) (`PyroOptim` method), 35
[scale](#) (`HalfCauchy` attribute), 26
[scale\(\)](#) (in module `pyro.poutine`), 41
[scale_tril](#) (`LowRankMultivariateNormal` attribute), 26

ScaleMessenger (class in `pyro.poutine.scale_messenger`),
47

`score_parts()` (Distribution method), 21

`score_parts()` (Rejector method), 27

`set_constraint()` (Parameterized method), 76

`set_data()` (GPModel method), 56

`set_mode()` (Parameterized method), 76

`set_prior()` (Parameterized method), 76

`set_state()` (ParamStoreDict method), 32

`set_state()` (PyroOptim method), 36

`setup()` (HMC method), 14

SGD() (in module `pyro.optim.pytorch_optimizers`), 36

`site_is_subsample()` (in module `pyro.poutine.util`), 49

SparseAdam() (in module
`pyro.optim.pytorch_optimizers`), 36

SparseGPRegression (class in
`pyro.contrib.gp.models.sgpr`), 59

`step()` (SVI method), 9

`stochastic_nodes` (Trace attribute), 44

StudentT (class in `pyro.distributions`), 20

`successors` (Trace attribute), 44

Sum (class in `pyro.contrib.gp.kernels.kernel`), 66

`support` (Binomial attribute), 24

`support` (Delta attribute), 24

`support` (Empirical attribute), 25

`support` (HalfCauchy attribute), 26

`support` (LowRankMultivariateNormal attribute), 26

`support` (VonMises attribute), 28

SVI (class in `pyro.infer.svi`), 9

T

TorchDistribution (class in `pyro.distributions`), 22

Trace (class in `pyro.poutine`), 42

`trace` (TraceHandler attribute), 47

`trace()` (in module `pyro.poutine`), 42

Trace_ELBO (class in `pyro.infer.trace_elbo`), 10

TraceEnum_ELBO (class in `pyro.infer.traceenum_elbo`),
11

TraceGraph_ELBO (class in `pyro.infer.tracegraph_elbo`),
11

TraceHandler (class in `pyro.poutine.trace_messenger`), 47

TraceMessenger (class in `pyro.poutine.trace_messenger`),
47

TracePosterior (class in `pyro.infer.abstract_infer`), 12

TracePredictive (class in `pyro.infer.abstract_infer`), 12

TransformedDistribution (class in `pyro.distributions`), 20

Transforming (class in `pyro.contrib.gp.kernels.kernel`), 66

U

Uniform (class in `pyro.distributions`), 20

`user_param_name()` (in module
`pyro.params.param_store`), 32

V

`validate_message()` (in module `pyro.poutine.runtime`), 48

`validation_enabled()` (in module `pyro`), 8

`variance` (Binomial attribute), 24

`variance` (Delta attribute), 24

`variance` (Empirical attribute), 25

`variance` (LowRankMultivariateNormal attribute), 26

VariationalGP (class in `pyro.contrib.gp.models.vgp`), 61

VariationalSparseGP (class in
`pyro.contrib.gp.models.vsgp`), 62

`vectorized` (CondIndepStackFrame attribute), 46

`vertical_scale()` (Kernel method), 65

VerticalScaling (class in `pyro.contrib.gp.kernels.kernel`),
67

VonMises (class in `pyro.distributions`), 28

W

`warp()` (Kernel method), 65

Warping (class in `pyro.contrib.gp.kernels.kernel`), 67

WhiteNoise (class in `pyro.contrib.gp.kernels.static`), 71