
Haiku

Haiku Contributors

Nov 30, 2023

BASICS

1	Installation	3
2	Known issues	173
3	Contribute	175
4	Support	177
5	License	179
6	Indices and tables	181
	Bibliography	183
	Python Module Index	185
	Index	187

Haiku is a library built on top of JAX designed to provide simple, composable abstractions for machine learning research.

```
import haiku as hk
import jax
import jax.numpy as jnp

def forward(x):
    mlp = hk.nets.MLP([300, 100, 10])
    return mlp(x)

forward = hk.transform(forward)

rng = hk.PRNGSequence(jax.random.PRNGKey(42))
x = jnp.ones([8, 28 * 28])
params = forward.init(next(rng), x)
logits = forward.apply(params, next(rng), x)
```


INSTALLATION

See <https://github.com/google/jax#pip-installation> for instructions on installing JAX.

We suggest installing the latest version of Haiku by running:

```
$ pip install git+https://github.com/deepmind/dm-haiku
```

Alternatively, you can install via PyPI:

```
$ pip install -U dm-haiku
```

1.1 Haiku Basics

In this Colab, you will learn the basics of Haiku.

What and Why ?

Haiku is a simple neural network library for JAX that enables users to use familiar object-oriented programming models while allowing full access to JAX’s pure function transformations. Haiku is designed to make the common things we do such as managing model parameters and other model state simpler and similar in spirit to the [Sonnet](#) library that has been widely used across DeepMind. It preserves Sonnet’s module-based programming model for state management while retaining access to JAX’s function transformations. Haiku can be expected to compose with other libraries and work well with the rest of JAX.

```
[2]: import haiku as hk
import jax
import jax.numpy as jnp
import numpy as np
```

1.1.1 A first example with `hk.transform`

As an initial introduction to Haiku, let us construct a linear module with weights and biases with custom initializations.

Similar to Sonnet modules, Haiku modules are Python objects that hold references to their own parameters, other modules, and methods that apply functions on user inputs. On the other hand, since JAX operates on pure function transformations, Haiku modules cannot be instantiated verbatim. Rather, the modules need to be wrapped into pure function transformations.

Haiku provides a simple function transformation, `hk.transform`, that turns functions that use these object-oriented, functionally “impure” modules into pure functions that can be used with JAX.

```
[3]: class MyLinear1(hk.Module):

    def __init__(self, output_size, name=None):
        super().__init__(name=name)
        self.output_size = output_size

    def __call__(self, x):
        j, k = x.shape[-1], self.output_size
        w_init = hk.initializers.TruncatedNormal(1. / np.sqrt(j))
        w = hk.get_parameter("w", shape=[j, k], dtype=x.dtype, init=w_init)
        b = hk.get_parameter("b", shape=[k], dtype=x.dtype, init=jnp.ones)
        return jnp.dot(x, w) + b
```

```
[4]: def _forward_fn_linear1(x):
    module = MyLinear1(output_size=2)
    return module(x)

forward_linear1 = hk.transform(_forward_fn_linear1)
```

We see that the forward wrapper object now contains two methods, `init` and `apply`, that are used to initialize the variables and do forward inference on the module.

```
[5]: forward_linear1
```

```
[5]: Transformed(init=<function without_state.<locals>.init_fn at 0x7f2e310354c0>, apply=
↳<function without_state.<locals>.apply_fn at 0x7f2e31035550>)
```

Calling the `init` method will initialize the parameters of the network and return them, as can be seen below. The `init` method takes a `jax.random.PRNGKey` and a sample input (usually just some dummy values to tell the networks about the expected shapes).

```
[6]: dummy_x = jnp.array([[1., 2., 3.]])
rng_key = jax.random.PRNGKey(42)

params = forward_linear1.init(rng=rng_key, x=dummy_x)
print(params)

{'my_linear1': {'w': DeviceArray([[[-0.30350363,  0.5123802 ],
 [ 0.08009142, -0.3163005 ]],
 [ 0.6056666 ,  0.5820702 ]]), dtype=float32), 'b': DeviceArray([1., 1.],
↳dtype=float32)}}
```

We can now use the `params` to apply the forward function to some inputs.

```
[7]: sample_x = jnp.array([[1., 2., 3.]])
sample_x_2 = jnp.array([[4., 5., 6.], [7., 8., 9.]])

output_1 = forward_linear1.apply(params=params, x=sample_x, rng=rng_key)
# Outputs are identical for given inputs since the forward inference is non-stochastic.
output_2 = forward_linear1.apply(params=params, x=sample_x, rng=rng_key)

output_3 = forward_linear1.apply(params=params, x=sample_x_2, rng=rng_key)

print(f'Output 1 : {output_1}')
```

(continues on next page)

(continued from previous page)

```
print(f'Output 2 (same as output 1): {output_2}')
print(f'Output 3 : {output_3}')
```

```
Output 1 : [[2.6736789 2.6259897]]
Output 2 (same as output 1): [[2.6736789 2.6259897]]
Output 3 : [[3.820442 4.960439]
            [4.967205 7.294889]]
```

Inference without random key

The module that we built is inherently non-stochastic. In that case, passing a random key to the apply method seems redundant. Haiku offers another transformation `hk.without_apply_rng` which can be further wrapped around our `hk.transform` method.

```
[8]: forward_without_rng = hk.without_apply_rng(hk.transform(_forward_fn_linear1))
      params = forward_without_rng.init(rng=rng_key, x=sample_x)
      output = forward_without_rng.apply(x=sample_x, params=params)
      print(f'Output without random key in forward pass \n {output_1}')
```

```
Output without random key in forward pass
[[2.6736789 2.6259897]]
```

We can also mutate the parameters and then do forward inference to generate a different output for the same inputs. This is what is done to apply gradient descent to our parameters while learning.

```
[9]: mutated_params = jax.tree_util.tree_map(lambda x: x+1., params)
      print(f'Mutated params \n : {mutated_params}')
      mutated_output = forward_without_rng.apply(x=sample_x, params=mutated_params)
      print(f'Output with mutated params \n {mutated_output}')
```

```
Mutated params
: {'my_linear1': {'b': DeviceArray([2., 2.], dtype=float32), 'w': DeviceArray([[0.
↪69649637, 1.5123801 ],
                [1.0800915 , 0.6836995 ],
                [1.6056666 , 1.5820701 ]], dtype=float32)}}
Output with mutated params
[[9.673679 9.62599 ]]
```

1.1.2 Stateful Inference in Haiku

For some modules you might want to maintain and carry over the internal state across function calls. Here, we demonstrate a simple example, where we declare a state variable `counter` within our Haiku transformation which gets updated on each call to the function. Note that we didn't explicitly instantiate this as a Haiku module (the same could be replicated as a `hk` module as shown earlier).

```
[10]: def stateful_f(x):
        counter = hk.get_state("counter", shape=[], dtype=jnp.int32, init=jnp.ones)
        multiplier = hk.get_parameter('multiplier', shape=[1,], dtype=x.dtype, init=jnp.ones)
        hk.set_state("counter", counter + 1)
        output = x + multiplier * counter
        return output

stateful_forward = hk.without_apply_rng(hk.transform_with_state(stateful_f))
```

(continues on next page)

(continued from previous page)

```

sample_x = jnp.array([[5., ]])
params, state = stateful_forward.init(x=sample_x, rng=rng_key)
print(f'Initial params:\n{params}\nInitial state:\n{state}')
print('#####')
for i in range(3):
    output, state = stateful_forward.apply(params, state, x=sample_x)
    print(f'After {i+1} iterations:\nOutput: {output}\nState: {state}')
    print('#####')

```

```

Initial params:
{'~': {'multiplier': DeviceArray([1.], dtype=float32)}}
Initial state:
{'~': {'counter': DeviceArray(1, dtype=int32)}}
#####
After 1 iterations:
Output: [[6.]]
State: {'~': {'counter': DeviceArray(2, dtype=int32)}}
#####
After 2 iterations:
Output: [[7.]]
State: {'~': {'counter': DeviceArray(3, dtype=int32)}}
#####
After 3 iterations:
Output: [[8.]]
State: {'~': {'counter': DeviceArray(4, dtype=int32)}}
#####

```

1.1.3 Built-in Haiku nets and nested modules

The usual networks we use such as MLP, Convnets etc. are defined already in Haiku and we can compose those modules to construct our custom Haiku Module.

Look at the params dictionary to see how the params are nested in the same way as the modules are nested within our custom Haiku module.

```

[11]: # See: https://dm-haiku.readthedocs.io/en/latest/api.html#common-modules

class MyModuleCustom(hk.Module):
    def __init__(self, output_size=2, name='custom_linear'):
        super().__init__(name=name)
        self._internal_linear_1 = hk.nets.MLP(output_sizes=[2, 3], name='hk_internal_linear')
        self._internal_linear_2 = MyLinear1(output_size=output_size, name='old_linear')

    def __call__(self, x):
        return self._internal_linear_2(self._internal_linear_1(x))

def _custom_forward_fn(x):
    module = MyModuleCustom()
    return module(x)

custom_forward_without_rng = hk.without_apply_rng(hk.transform(_custom_forward_fn))

```

(continues on next page)

(continued from previous page)

```

params = custom_forward_without_rng.init(rng=rng_key, x=sample_x)
params
[11]: {'custom_linear~/hk_internal_linear~/~/linear_0': {'b': DeviceArray([0., 0.],
↳ dtype=float32),
      'w': DeviceArray([[ 1.51595   , -0.23353337]], dtype=float32)},
      'custom_linear~/hk_internal_linear~/~/linear_1': {'b': DeviceArray([0., 0., 0.],
↳ dtype=float32),
      'w': DeviceArray([[ -0.22075887, -0.27375957,  0.5931483 ],
                        [ 0.7818068 ,  0.72626334, -0.6860752 ]], dtype=float32)},
      'custom_linear~/old_linear': {'b': DeviceArray([1., 1.], dtype=float32),
      'w': DeviceArray([[ 0.28584382,  0.31626168],
                        [ 0.2335775 , -0.4827032 ],
                        [-0.14647584, -0.7185701 ]], dtype=float32)}}

```

1.1.4 RNG Keys with `hk.next_rng_key()`

The modules that we saw earlier were all non-stochastic. Below we show how to sample random numbers to do stochastic inference.

Haiku offers a trivial model for working with random numbers. Within a transformed function, `hk.next_rng_key()` returns a unique rng key. These unique keys are deterministically derived from an initial random key passed into the top-level transformed function, and are thus safe to use with JAX program transformations.

Let us define a simple haiku function where we generate two random samples. Note that the `next_rng_keys` are determined from the initial random key passed to the `apply` method of the top-level transformed function.

```

[15]: class HkRandom2(hk.Module):
      def __init__(self, rate=0.5):
          super().__init__()
          self.rate = rate

      def __call__(self, x):
          key1 = hk.next_rng_key()
          return jax.random.bernoulli(key1, 1.0 - self.rate, shape=x.shape)

class HkRandomNest(hk.Module):
      def __init__(self, rate=0.5):
          super().__init__()
          self.rate = rate
          self._another_random_module = HkRandom2()

      def __call__(self, x):
          key2 = hk.next_rng_key()
          p1 = self._another_random_module(x)
          p2 = jax.random.bernoulli(key2, 1.0 - self.rate, shape=x.shape)
          print(f'Bernoullis are : {p1, p2}')

# Note that the modules that are stochastic cannot be wrapped with hk.without_apply_rng()
forward = hk.transform(lambda x: HkRandomNest()(x))

```

(continues on next page)

(continued from previous page)

```
x = jnp.array(1.)
print("INIT:")
params = forward.init(rng_key, x=x)
print("APPLY:")
prediction = forward.apply(params, x=x, rng=rng_key)
```

```
INIT:
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
APPLY:
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
```

Note that this means that passing the same random key to multiple calls of the apply function will generate the same stochastic results!

```
[16]: for _ in range(3):
        forward.apply(params, x=x, rng=rng_key)
```

```
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
```

Make sure to split off new RNG keys to get different stochastic behavior across apply calls, and to never re-use an RNG key. (For a more comprehensive explanation of how to handle random state in JAX, check out this RNG tutorial: <https://jax.readthedocs.io/en/latest/jax-101/05-random-numbers.html>.)

```
[19]: for _ in range(3):
        rng_key, apply_rng_key = jax.random.split(rng_key)
        forward.apply(params, x=x, rng=apply_rng_key)
```

```
Bernoullis are : (DeviceArray(False, dtype=bool), DeviceArray(False, dtype=bool))
Bernoullis are : (DeviceArray(False, dtype=bool), DeviceArray(True, dtype=bool))
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(False, dtype=bool))
```

Haiku also provides `hk.PRNGSequence` which returns an iterator of random keys.

```
[20]: rng_sequence = hk.PRNGSequence(rng_key)
        for _ in range(3):
            forward.apply(params, x=x, rng=next(rng_sequence))
```

```
Bernoullis are : (DeviceArray(True, dtype=bool), DeviceArray(True, dtype=bool))
Bernoullis are : (DeviceArray(False, dtype=bool), DeviceArray(False, dtype=bool))
Bernoullis are : (DeviceArray(False, dtype=bool), DeviceArray(True, dtype=bool))
```

```
[ ]: import haiku as hk
        import jax
        import jax.numpy as jnp
```

TL;DR: A JAX transform inside of a `hk.transform` is likely to transform a side effecting function, which will result in an `UnexpectedTracerError`. This page describes two ways to get around this.

1.2 Limitations of Nesting JAX Functions and Haiku Modules

Once a Haiku network has been transformed to a pair of pure functions using `hk.transform`, it's possible to freely combine these with any JAX transformations like `jax.jit`, `jax.grad`, `jax.lax.scan` and so on.

If you want to use JAX transformations **inside** of a `hk.transform` however, you need to be more careful. It's possible, but most functions inside of the `hk.transform` boundary are still side effecting, and cannot safely be transformed by JAX. This is a common cause of `UnexpectedTracerErrors` in code using Haiku. These errors are a result of using a JAX transform on a side effecting function (for more information on this JAX error, see <https://jax.readthedocs.io/en/latest/errors.html#jax.errors.UnexpectedTracerError>).

An example with `jax.eval_shape`:

```
[ ]: def net(x): # inside of a hk.transform, this is still side-effecting
      w = hk.get_parameter("w", (2, 2), init=jnp.ones)
      return w @ x

def eval_shape_net(x):
    output_shape = jax.eval_shape(net, x) # eval_shape on side-effecting function
    return net(x)                          # UnexpectedTracerError!

init, _ = hk.transform(eval_shape_net)
try:
    init(jax.random.PRNGKey(666), jnp.ones((2, 2)))
except jax.errors.UnexpectedTracerError:
    print("UnexpectedTracerError: applied JAX transform to side effecting function")

WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and rerun for more info.)

UnexpectedTracerError: applied JAX transform to side effecting function
```

These examples use `jax.eval_shape`, but could have used any higher-order JAX function (eg. `jax.vmap`, `jax.lax.scan`, `jax.while_loop`, ...).

The error points to `hk.get_parameter`. This is the operation which makes `net` a side effecting function. The side effect in this case is the creation of a parameter, which gets stored into the Haiku state. Similarly you would get an error using `hk.next_rng_key`, because it advances the Haiku RNG state and stores a new `PRNGKey` into the Haiku state. In general, transforming a non-transformed Haiku module will result in an `UnexpectedTracerError`.

You could re-write the code above to create the parameter outside of the `eval_shape` transformation, making `net` a pure function by threading through the parameter explicitly as an argument:

```
[ ]: def net(w, x): # no side effects!
      return w @ x

def eval_shape_net(x):
    w = hk.get_parameter("w", (3, 2), init=jnp.ones)
    output_shape = jax.eval_shape(net, w, x) # net is now side-effect free
    return output_shape, net(w, x)

key = jax.random.PRNGKey(777)
x = jnp.ones((2, 3))
init, apply = hk.transform(eval_shape_net)
params = init(key, x)
apply(params, key, x)
```

```
(ShapeDtypeStruct(shape=(3, 3), dtype=float32),
 DeviceArray([[2., 2., 2.],
              [2., 2., 2.],
              [2., 2., 2.]], dtype=float32))
```

However, that's not always possible. Consider the following code which calls a Haiku module (`hk.nets.MLP`) which we don't own. This module will internally call `get_parameter`.

```
[ ]: def eval_shape_net(x):
    net = hk.nets.MLP([300, 100])
    output_shape = jax.eval_shape(net, x)
    return output_shape, net(x)

init, _ = hk.transform(eval_shape_net)
try:
    init(jax.random.PRNGKey(666), jnp.ones((2, 2)))
except jax.errors.UnexpectedTracerError:
    print("UnexpectedTracerError: applied JAX transform to side effecting function")
```

```
UnexpectedTracerError: applied JAX transform to side effecting function
```

1.2.1 Using `hk.lift`

We want a way to get access to our implicit Haiku state, and get a functionally pure version of `hk.nets.MLP`. The way to usually achieve this is by using a `hk.transform`, so all we need is a way to nest an inner `hk.transform` inside an outer `hk.transform`! We'll create another pair of `init` and `apply` functions through `hk.transform`, and these can then be safely combined with any higher-order JAX function.

However, we need a way to register this nested `hk.transform` state into the outer scope. We can use `hk.lift` for this. Wrapping our inner `init` function with `hk.lift` will register our inner params into the outer parameter scope.

```
[ ]: def eval_shape_net(x):
    net = hk.nets.MLP([300, 100]) # still side-effecting
    init, apply = hk.transform(net) # nested transform
    params = hk.lift(init, name="inner")(hk.next_rng_key(), x) # register parameters in
    ↪outer module scope with name "inner"
    output_shape = jax.eval_shape(apply, params, hk.next_rng_key(), x) # apply is a
    ↪functionally pure function and can be transformed!
    out = net(x)
    return out, output_shape
```

```
init, apply = hk.transform(eval_shape_net)
params = init(jax.random.PRNGKey(777), jnp.ones((100, 100)))
apply(params, jax.random.PRNGKey(777), jnp.ones((100, 100)))
jax.tree_util.tree_map(lambda x: x.shape, params)
```

```
FlatMap({
  'inner/mlp/~linear_0': FlatMap({'b': (300,), 'w': (100, 300)}),
  'inner/mlp/~linear_1': FlatMap({'b': (100,), 'w': (300, 100)}),
  'mlp/~linear_0': FlatMap({'b': (300,), 'w': (100, 300)}),
  'mlp/~linear_1': FlatMap({'b': (100,), 'w': (300, 100)}),
})
```

1.2.2 Using Haiku versions of JAX transforms

Haiku also provides wrapped versions of some of the JAX functions for convenience. For example: `hk.grad`, `hk.vmap`, ... See <https://dm-haiku.readthedocs.io/en/latest/api.html#jax-fundamentals> for a full list of available functions.

These wrappers apply the JAX function to a functionally pure version of the Haiku function, by doing the explicit state threading for you. They don't introduce an extra namescoping level like `lift` does.

```
[ ]: def eval_shape_net(x):
    net = hk.nets.MLP([300, 100])          # still side-effecting
    output_shape = hk.eval_shape(net, x)  # hk.eval_shape threads through the Haiku state,
    ↪ for you
    out = net(x)
    return out, output_shape

init, apply = hk.transform(eval_shape_net)
params = init(jax.random.PRNGKey(777), jnp.ones((100, 100)))
out = apply(params, jax.random.PRNGKey(777), jnp.ones((100, 100)))
```

1.2.3 Summary

To summarize, some good and bad examples of combining JAX transforms and Haiku modules:

What?	Works?	Example
vmapping outside a <code>hk.transform</code>	✓ yes!	<code>jax.vmap(hk.transform(hk.nets.ResNet50))</code>
vmapping inside a <code>hk.transform</code>	✗ unexpected tracer error	<code>hk.transform(jax.vmap(hk.nets.ResNet50))</code>
vmapping a nested <code>hk.transform</code> (without <code>lift</code>)	✗ inner state is not registered	<code>hk.transform(jax.vmap(hk.transform(hk.nets.ResNet50)))</code>
vmapping a nested <code>hk.transform</code> (with <code>lift</code>)	✓ yes!	<code>hk.transform(jax.vmap(hk.lift(hk.transform(hk.nets.ResNet50))))</code>
using <code>hk.vmap</code>	✓ yes!	<code>hk.transform(hk.vmap(hk.nets.ResNet50))</code>

1.3 Haiku API reference

1.3.1 Haiku Fundamentals

Haiku Transforms

<code>transform(f, *, [apply_rng])</code>	Transforms a function using Haiku modules into a pair of pure functions.
<code>transform_with_state(f)</code>	Transforms a function using Haiku modules into a pair of pure functions.
<code>multi_transform(f)</code>	Transforms a collection of functions using Haiku into pure functions.
<code>multi_transform_with_state(f)</code>	Transforms a collection of functions using Haiku into pure functions.
<code>without_apply_rng(f)</code>	Removes the <code>rng</code> argument from the <code>apply</code> function.
<code>without_state(f)</code>	Wraps a transformed tuple and ignores state in/out.

transform

`haiku.transform(f, *, apply_rng=True)`

Transforms a function using Haiku modules into a pair of pure functions.

For a function `out = f(*a, **k)` this function returns a pair of two pure functions that call `f(*a, **k)` explicitly collecting and injecting parameter values:

```
params = init(rng, *a, **k)
out = apply(params, rng, *a, **k)
```

Note that the `rng` argument is typically not required for `apply` and passing `None` is accepted.

The first thing to do is to define a *Module*. A module encapsulates some parameters and a computation on those parameters:

```
>>> class MyModule(hk.Module):
...     def __call__(self, x):
...         w = hk.get_parameter("w", [], init=jnp.zeros)
...         return x + w
```

Next, define some function that creates and applies modules. We use `transform()` to transform that function into a pair of functions that allow us to lift all the parameters out of the function (`f.init`) and apply the function with a given set of parameters (`f.apply`):

```
>>> def f(x):
...     a = MyModule()
...     b = MyModule()
...     return a(x) + b(x)
>>> f = hk.transform(f)
```

To get the initial state of the module call `init` with an example input:

```
>>> params = f.init(None, 1)
>>> params
{'my_module': {'w': ...Array(0., dtype=float32)},
 'my_module_1': {'w': ...Array(0., dtype=float32)}}
```

You can then apply the function with the given parameters by calling `apply` (note that since we don't use Haiku's random number APIs to apply our network we pass `None` as an RNG key):

```
>>> print(f.apply(params, None, 1))
2.0
```

It is expected that your program will at some point produce updated parameters and you will want to re-apply `apply`. You can do this by calling `apply` with different parameters:

```
>>> new_params = {"my_module": {"w": jnp.array(2.)},
...               "my_module_1": {"w": jnp.array(3.)}}
>>> print(f.apply(new_params, None, 2))
9.0
```

If your transformed function needs to maintain internal state (e.g. moving averages in batch norm) then see `transform_with_state()`.

Parameters

- **f** – A function closing over *Module* instances.
- **apply_rng** – In the process of being removed. Can only value *True*.

Return type *Transformed*

Returns A *Transformed* tuple with `init` and `apply` pure functions.

transform_with_state

`haiku.transform_with_state(f)`

Transforms a function using Haiku modules into a pair of pure functions.

See `transform()` for general details on Haiku transformations.

For a function `out = f(*a, **k)` this function returns a pair of two pure functions that call `f(*a, **k)` explicitly collecting and injecting parameter values and state:

```
params, state = init(rng, *a, **k)
out, state = apply(params, state, rng, *a, **k)
```

Note that the `rng` argument is typically not required for `apply` and passing `None` is accepted.

This function is equivalent to `transform()`, however it allows you to maintain and update internal state (e.g. *ExponentialMovingAverage* in *BatchNorm*) via `get_state()` and `set_state()`:

```
>>> def f():
...     counter = hk.get_state("counter", shape=[], dtype=jnp.int32,
...                             init=jnp.zeros)
...     hk.set_state("counter", counter + 1)
...     return counter
>>> f = hk.transform_with_state(f)
```

```
>>> params, state = f.init(None)
>>> for _ in range(10):
...     counter, state = f.apply(params, state, None)
>>> print(counter)
9
```

Parameters **f** – A function closing over *Module* instances.

Return type *TransformedWithState*

Returns A *TransformedWithState* tuple with `init` and `apply` pure functions.

multi_transform

`haiku.multi_transform(f)`

Transforms a collection of functions using Haiku into pure functions.

In many scenarios we have several modules which are used either as primitives for several Haiku modules/functions, or whose pure versions are to be reused in downstream code. This utility enables this by applying `transform()` to an arbitrary tree of Haiku functions which share modules and have a common `init` function.

`f` is expected to return a tuple of two elements. First is a `template` Haiku function which provides an example of how all internal Haiku modules are connected. This function is used to create a common `init` function (with your parameters).

The second object is an arbitrary tree of Haiku functions all of which reuse the modules connected in the `template` function. These functions are transformed to pure `apply` functions.

Example:

```
>>> def f():
...     encoder = hk.Linear(1, name="encoder")
...     decoder = hk.Linear(1, name="decoder")
...
...     def init(x):
...         z = encoder(x)
...         return decoder(z)
...
...     return init, (encoder, decoder)
```

```
>>> f = hk.multi_transform(f)
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 1])
>>> params = f.init(rng, x)
>>> jax.tree_util.tree_map(jnp.shape, params)
{'decoder': {'b': (1,), 'w': (1, 1)},
 'encoder': {'b': (1,), 'w': (1, 1)}}
```

```
>>> encode, decode = f.apply
>>> z = encode(params, None, x)
>>> y = decode(params, None, z)
```

Parameters `f` (`Callable[[], tuple[TemplateFn, TreeOfApplyFns]]`) – A factory function that returns two functions, firstly a common `init` function that creates all modules, and secondly a pytree of `apply` functions which make use of those modules.

Return type `MultiTransformed`

Returns

A `MultiTransformed` instance which contains a pure `init` function that creates all parameters, and a pytree of pure `apply` functions that given the params apply the given function.

See also:

`multi_transform_with_state()`: Equivalent for modules using state.

multi_transform_with_state

`haiku.multi_transform_with_state(f)`

Transforms a collection of functions using Haiku into pure functions.

See `multi_transform()` for more details.

Example:

```
>>> def f():
...     encoder = hk.Linear(1, name="encoder")
...     decoder = hk.Linear(1, name="decoder")
...
...     def init(x):
...         z = encoder(x)
...         return decoder(z)
...
...     return init, (encoder, decoder)
```

```
>>> f = hk.multi_transform_with_state(f)
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 1])
>>> params, state = f.init(rng, x)
>>> jax.tree_util.tree_map(jnp.shape, params)
{'decoder': {'b': (1,), 'w': (1, 1)},
 'encoder': {'b': (1,), 'w': (1, 1)}}
```

```
>>> encode, decode = f.apply
>>> z, state = encode(params, state, None, x)
>>> y, state = decode(params, state, None, z)
```

Parameters `f` (`Callable[[], tuple[TemplateFn, TreeOfApplyFns]]`) – Function returning a “template” function and an arbitrary tree of functions using modules connected in the template function.

Return type `MultiTransformedWithState`

Returns An `init` function and a tree of pure apply functions.

See also:

- `transform_with_state()`: Transform a single apply function.
- `multi_transform()`: Transform multiple apply functions without state.

without_apply_rng

`haiku.without_apply_rng(f)`

Removes the `rng` argument from the apply function.

This is a convenience wrapper that makes the `rng` argument to `f.apply` default to `None`. This is useful when `f` doesn't actually use random numbers as part of its computation, such that the `rng` argument wouldn't be used. Note that if `f` *does* use random numbers, this will cause an error to be thrown complaining that `f` needs a non-`None` `PRNGKey`.

Parameters `f` (*TransformedT*) – A transformed function.

Return type *TransformedT*

Returns The same transformed function, with a modified `apply`.

without_state

`haiku.without_state(f)`

Wraps a transformed tuple and ignores state in/out.

The example below is equivalent to `f = hk.transform(f)`:

```

>>> def f(x):
...     mod = hk.Linear(10)
...     return mod(x)
>>> f = hk.without_state(hk.transform_with_state(f))
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.zeros([1, 1])
>>> params = f.init(rng, x)
>>> print(f.apply(params, rng, x))
[[0. 0. 0. 0. 0. 0. 0. 0. 0. 0.]]

```

Parameters `f` (*TransformedWithState*) – A transformed function.

Return type *Transformed*

Returns A transformed function that does not take or return state.

with_empty_state

`haiku.with_empty_state(f)`

Wraps a transformed tuple and passes empty state in/out.

The example below is equivalent to `f = hk.transform_with_state(f)`:

```

>>> def f(x):
...     mod = hk.Linear(10)
...     return mod(x)
>>> f = hk.with_empty_state(hk.transform(f))
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.zeros([1, 1])
>>> params, state = f.init(rng, x)
>>> state

```

(continues on next page)

(continued from previous page)

```

{}
>>> out, state = f.apply(params, state, rng, x)
>>> print(out)
[[0. 0. 0. 0. 0. 0. 0. 0. 0. 0.]]
>>> state
{}

```

Parameters *f* (*Transformed*) – A transformed function.

Return type *TransformedWithState*

Returns A transformed function that does accepts and returns state.

Modules, Parameters and State

<i>Module</i> ([name])	Base class for Haiku modules.
<i>to_module</i> (f)	Converts a function into a callable module class.
<i>get_parameter</i> (name, shape[, dtype, init])	Creates or reuses a parameter for the given transformed function.
<i>get_state</i> (name[, shape, dtype, init])	Gets the current value for state with an optional initializer.
<i>set_state</i> (name, value)	Sets the current value for some state.

Module

class `haiku.Module`(*name=None*)

Base class for Haiku modules.

A Haiku module is a lightweight container for variables and other modules. Modules typically define one or more “forward” methods (e.g. `__call__`) which apply operations combining user input and module parameters.

Modules must be initialized inside a `transform()` call.

For example:

```

>>> class AddModule(hk.Module):
...     def __call__(self, x):
...         w = hk.get_parameter("w", [], init=jnp.ones)
...         return x + w

```

```

>>> def forward_fn(x):
...     mod = AddModule()
...     return mod(x)

```

```

>>> forward = hk.transform(forward_fn)
>>> x = 1.
>>> rng = None
>>> params = forward.init(rng, x)
>>> print(forward.apply(params, None, x))
2.0

```

`__init__(name=None)`

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters `name` (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If `name` is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

`__post_init__(name=None)`

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters `name` (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If `name` is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

`params_dict()`

Returns parameters keyed by name for this module and submodules.

Return type Mapping[str, jnp.ndarray]

`state_dict()`

Returns state keyed by name for this module and submodules.

Return type Mapping[str, jnp.ndarray]

to_module

`haiku.to_module(f)`

Converts a function into a callable module class.

Sample usage:

```
>>> def bias_fn(x):
...     b = hk.get_parameter("b", [], init=hk.initializers.RandomNormal())
...     return x + b
>>> Bias = hk.to_module(bias_fn)
>>> def net(x, y):
...     b = Bias(name="my_bias")
...     # Bias x and y by the same amount.
...     return b(x) * b(y)
```

Parameters `f` (*Callable[... , Any]*) – The function to convert.

Return type type[CallableModule]

Returns A module class which runs `f` when called.

get_parameter

`haiku.get_parameter(name, shape, dtype=<class 'jax.numpy.float32'>, init=None)`

Creates or reuses a parameter for the given transformed function.

```
>>> print(hk.get_parameter("w", [], init=jnp.ones))
1.0
```

Parameters within the same `transform()` and/or `Module` with the same name have the same value:

```
>>> w1 = hk.get_parameter("w", [], init=jnp.zeros)
>>> w2 = hk.get_parameter("w", [], init=jnp.zeros)
>>> assert w1 is w2
```

Parameters

- **name** (*str*) – A name for the parameter.
- **shape** (*Sequence[int]*) – The shape of the parameter.
- **dtype** (*Any*) – The dtype of the parameter.
- **init** (*Optional[Initializer]*) – A callable of shape, dtype to generate an initial value for the parameter.

Return type `jax.Array`

Returns A `jax.Array` with the parameter of the given shape.

get_state

`haiku.get_state(name, shape=None, dtype=<class 'jax.numpy.float32'>, init=None)`

Gets the current value for state with an optional initializer.

“State” can be used to represent mutable state in your network. The most common usage of state is to represent the moving averages used in batch normalization (see [ExponentialMovingAverage](#)). If your network uses “state” then you are required to use `transform_with_state()` and pass state into and out of the apply function.

```
>>> print(hk.get_state("counter", [], init=jnp.zeros))
0.0
```

If the value for the given state is already defined (e.g. using `set_state()`) then you can call with just the name:

```
>>> print(hk.get_state("counter"))
0.0
```

NOTE: state within the same `transform()` and/or `Module` with the same name have the same value:

```
>>> c1 = hk.get_state("counter")
>>> c2 = hk.get_state("counter")
>>> assert c1 is c2
```

Parameters

- **name** (*str*) – A name for the state.
- **shape** (*Optional[Sequence[int]]*) – The shape of the state.

- **dtype** (*Any*) – The dtype of the state.
- **init** (*Optional[Initializer]*) – A callable $f(\text{shape}, \text{dtype})$ that returns an initial value for the state.

Return type `jax.Array`

Returns A `jax.Array` with the state of the given shape.

set_state

`haiku.set_state(name, value)`

Sets the current value for some state.

See `get_state()`.

“State” can be used to represent mutable state in your network. The most common usage of state is to represent the moving averages used in batch normalization (see [ExponentialMovingAverage](#)). If your network uses “state” then you are required to use `transform_with_state()` and pass state into and out of the apply function.

```
>>> hk.set_state("counter", jnp.zeros([]])
>>> print(hk.get_state("counter"))
0.0
```

NOTE: state within the same `transform()` and/or `Module` with the same name have the same value:

```
>>> w1 = hk.get_state("counter")
>>> w2 = hk.get_state("counter")
>>> assert w1 is w2
```

Parameters

- **name** (*str*) – A name for the state.
- **value** – A value to set.

Getters and Interceptors

<code>custom_creator(creator, *[, params, state])</code>	Registers a custom parameter and/or state creator.
<code>custom_getter(getter, *[, params, state])</code>	Registers a custom parameter or state getter.
<code>custom_setter(setter)</code>	Registers a custom state setter.
<code>GetterContext(full_name, module, ...)</code>	Context about where parameters are being created.
<code>SetterContext(full_name, module, ...)</code>	Context about where state is being set.
<code>intercept_methods(interceptor)</code>	Register a new method interceptor.
<code>MethodContext(module, method_name, ...)</code>	Read only state showing the calling context for a method.

custom_creator

`haiku.custom_creator(creator, *, params=True, state=False)`

Registers a custom parameter and/or state creator.

When new parameters are created via `get_parameter()` we first run custom creators passing user defined values through. For example:

```
>>> def zeros_creator(next_creator, shape, dtype, init, context):
...     init = jnp.zeros
...     return next_creator(shape, dtype, init)
```

```
>>> with hk.custom_creator(zeros_creator):
...     z = hk.get_parameter("z", [], jnp.float32, jnp.ones)
>>> print(z)
0.0
```

If `state=True` then your creator will additionally run on calls to `get_state()`:

```
>>> with hk.custom_creator(zeros_creator, state=True):
...     z = hk.get_state("z", [], jnp.float32, jnp.ones)
>>> print(z)
0.0
```

Parameters

- **creator** (*Creator*) – A parameter creator.
- **params** (*bool*) – Whether to intercept parameter creation, defaults to `True`.
- **state** (*bool*) – Whether to intercept state creation, defaults to `False`.

Return type `contextlib.AbstractContextManager`

Returns Context manager under which the creator is active.

custom_getter

`haiku.custom_getter(getter, *, params=True, state=False)`

Registers a custom parameter or state getter.

When parameters are retrieved using `get_parameter()` we always run all custom getters before returning a value to the user.

```
>>> def bf16_getter(next_getter, value, context):
...     value = value.astype(jnp.bfloat16)
...     return next_getter(value)
```

```
>>> with hk.custom_getter(bf16_getter):
...     w = hk.get_parameter("w", [], jnp.float32, jnp.ones)
>>> w.dtype
dtype(bfloat16)
```

If `state=True` the getter will additionally run for calls to `get_state()`:

```
>>> with hk.custom_getter(bf16_getter, state=True):
...     c = hk.get_state("c", [], jnp.float32, jnp.ones)
>>> c.dtype
dtype(bfloat16)
```

Parameters

- **getter** (*Getter*) – A parameter getter.
- **params** (*bool*) – Whether the getter should run on `get_parameter()`
- **state** (*bool*) – Whether the getter should run on `get_state()`.

Return type `contextlib.AbstractContextManager`

Returns Context manager under which the getter is active.

custom_setter

`haiku.custom_setter(setter)`

Registers a custom state setter.

When state is set using `set_state()` we always run all custom setters before saving the value.

```
>>> def zero_during_init(next_setter, value, context):
...     if hk.running_init():
...         value = jnp.zeros_like(value)
...     return next_setter(value)
```

```
>>> with hk.custom_setter(zero_during_init):
...     hk.set_state("x", jnp.ones([2]))
...     x = hk.get_state("x")
>>> print(x)
[0. 0.]
```

Parameters **setter** (*Setter*) – A state setter.

Return type `contextlib.AbstractContextManager`

Returns Context manager under which the setter is active.

GetterContext

```
class haiku.GetterContext(full_name: str, module: Optional[Module], original_dtype: Any, original_shape:
    Sequence[int], original_init: Optional[Initializer], lifted_prefix_name:
    Optional[str])
```

Context about where parameters are being created.

full_name

The full name of the given parameter (e.g. `mlp/~/linear_0/w`).

Type `str`

module

The module that owns the current parameter, `None` if this parameter exists outside any module.

Type `Optional[Module]`

original_dtype

The dtype that `get_parameter()` or `get_state()` was originally called with.

Type `Any`

original_shape

The shape that `get_parameter()` or `get_state()` was originally called with.

Type `Sequence[int]`

original_init

The initializer that `get_parameter()` or `get_state()` was originally called with.

Type `Optional[Initializer]`

lifted_prefix_name

The module names of all enclosing lifted modules (see `lift()` for more context). Adding this string as a prefix to `full_name` will be equal to the final parameter name in the outer transform's parameter dictionary. NOTE: When `get_parameter()` or `get_state()` is called in an `apply` context, this name will always be `None` because only `init` functions are lifted.

Type `Optional[str]`

module_name

The full name of enclosing modules.

name

The name of this parameter.

SetterContext

```
class haiku.SetterContext(full_name: str, module: Optional[Module], original_dtype: Any, original_shape:
    Sequence[int], lifted_prefix_name: Optional[str])
```

Context about where state is being set.

full_name

The full name of the given state (e.g. `mlp/~~/linear_0/w`).

Type `str`

module

The module that owns the current state, `None` if this state exists outside any module.

Type `Optional[Module]`

original_dtype

The dtype that `set_state()` was originally called with.

Type `Any`

original_shape

The shape that `set_state()` or `get_state()` was originally called with.

Type `Sequence[int]`

lifted_prefix_name

The module names of all enclosing lifted modules (see `lift()` for more context). Adding this string as a prefix to `full_name` will be equal to the final parameter name in the outer transform's parameter dictionary. NOTE: When `get_parameter()` or `get_state()` is called in an `apply` context, this name will always be None because only `init` functions are lifted.

Type Optional[str]

module_name

The full name of enclosing modules.

name

The name of this state.

intercept_methods

`haiku.intercept_methods(interceptor)`

Register a new method interceptor.

Method interceptors allow you to (at a distance) intercept method calls to modules and modify args/kwargs before calling the underlying method. After the underlying method is called you can modify its result before it is passed back to the user.

For example you could intercept method calls to `BatchNorm` and ensure it is always computed in full precision:

```
>>> def my_interceptor(next_f, args, kwargs, context):
...     if (type(context.module) is not hk.BatchNorm
...         or context.method_name != "__call__"):
...         # We ignore methods other than BatchNorm.__call__.
...         return next_f(*args, **kwargs)
...
...     def cast_if_array(x):
...         if isinstance(x, jax.Array):
...             x = x.astype(jnp.float32)
...         return x
...
...     args, kwargs = jax.tree_util.tree_map(cast_if_array, (args, kwargs))
...     out = next_f(*args, **kwargs)
...     return out
```

We can create and use our module in the usual way, we just need to wrap any method calls we want to intercept in the context manager:

```
>>> mod = hk.BatchNorm(decay_rate=0.9, create_scale=True, create_offset=True)
>>> x = jnp.ones([], jnp.bfloat16)
>>> with hk.intercept_methods(my_interceptor):
...     out = mod(x, is_training=True)
>>> assert out.dtype == jnp.float32
```

Without the interceptor `BatchNorm` would compute in bf16, however since we cast `x` before the underlying method is called we compute in f32.

Parameters `interceptor` (*MethodGetter*) – A method interceptor.

Returns Context manager under which the interceptor is active.

MethodContext

```
class haiku.MethodContext(module: 'Module', method_name: str, orig_method: Callable[..., Any], orig_class:
    type['Module'])
```

Read only state showing the calling context for a method.

For example, let's define two interceptors and print the values in the context. Additionally, we will make the first interceptor conditionally short circuit, since interceptors stack and are run in order, an earlier interceptor can decide to call the next interceptor, or short circuit and call the underlying method directly:

```
>>> module = hk.Linear(1, name="method_context_example")
>>> short_circuit = False

>>> def my_interceptor_1(next_fun, args, kwargs, context):
...     print('running my_interceptor_1')
...     print('- module.name: ', context.module.name)
...     print('- method_name: ', context.method_name)
...     if short_circuit:
...         return context.orig_method(*args, **kwargs)
...     else:
...         return next_fun(*args, **kwargs)
>>> def my_interceptor_2(next_fun, args, kwargs, context):
...     print('running my_interceptor_2')
...     print('- module.name: ', context.module.name)
...     print('- method_name: ', context.method_name)
...     return next_fun(*args, **kwargs)
```

When short_circuit=False the two interceptors will run in order:

```
>>> with hk.intercept_methods(my_interceptor_1), \
...     hk.intercept_methods(my_interceptor_2):
...     _ = module(jnp.ones([1, 1]))
running my_interceptor_1
- module.name: method_context_example
- method_name: __call__
running my_interceptor_2
- module.name: method_context_example
- method_name: __call__
```

Setting short_circuit=True will cause the first interceptor to call the original method (rather than next_fun which will trigger the next interceptor):

```
>>> short_circuit = True
>>> with hk.intercept_methods(my_interceptor_1), \
...     hk.intercept_methods(my_interceptor_2):
...     _ = module(jnp.ones([1, 1]))
running my_interceptor_1
- module.name: method_context_example
- method_name: __call__
```

module

A *Module* instance whose method is being called.

Type 'Module'

method_name

The name of the method being called on the module.

Type str

orig_method

The underlying method on the module which when called will *not* trigger interceptors. You should only call this if you want to short circuit all the other interceptors, in general you should prefer to call the `next_fun` passed to your interceptor which will run `orig_method` after running all other interceptors.

Type Callable[..., Any]

orig_class

The class which defined `orig_method`. Note that when using inheritance this is not necessarily the same as `type(module)`.

Type type['Module']

Random Numbers

<code>PRNGSequence(key_or_seed)</code>	Iterator of JAX random keys.
<code>next_rng_key()</code>	Returns a unique JAX random key split from the current global key.
<code>next_rng_keys(num)</code>	Returns one or more JAX random keys split from the current global key.
<code>maybe_next_rng_key()</code>	<code>next_rng_key()</code> if random numbers are available, else <code>None</code> .
<code>reserve_rng_keys(num)</code>	Pre-allocate some number of JAX RNG keys.
<code>with_rng(key)</code>	Provides a new sequence for <code>next_rng_key()</code> to draw from.
<code>maybe_get_rng_sequence_state()</code>	Returns the internal state of the PRNG sequence.
<code>replace_rng_sequence_state(state)</code>	Replaces the internal state of the PRNG sequence with the given state.

PRNGSequence

class haiku.PRNGSequence(*key_or_seed*)

Iterator of JAX random keys.

```
>>> seq = hk.PRNGSequence(42) # OR pass a jax.random.PRNGKey
>>> key1 = next(seq)
>>> key2 = next(seq)
>>> assert key1 is not key2
```

If you know how many keys you will want then you can use `reserve()` to more efficiently split the keys you need:

```
>>> seq.reserve(4)
>>> keys = [next(seq) for _ in range(4)]
```

__init__(*key_or_seed*)

Creates a new `PRNGSequence`.

reserve(*num*)

Splits additional *num* keys for later use.

__next__()

Return the next item from the iterator. When exhausted, raise `StopIteration`

Return type PRNGKey

next()

Return the next item from the iterator. When exhausted, raise `StopIteration`

Return type PRNGKey

next_rng_key**haiku.next_rng_key**()

Returns a unique JAX random key split from the current global key.

```
>>> key = hk.next_rng_key()
>>> _ = jax.random.uniform(key, [])
```

Return type PRNGKey

Returns A unique (within a call to `init` or `apply`) JAX rng key that can be used with APIs such as `jax.random.uniform()`.

next_rng_keys**haiku.next_rng_keys**(*num*)

Returns one or more JAX random keys split from the current global key.

```
>>> k1, k2 = hk.next_rng_keys(2)
>>> assert (k1 != k2).all()
>>> a = jax.random.uniform(k1, [])
>>> b = jax.random.uniform(k2, [])
>>> assert a != b
```

Parameters *num* (*int*) – The number of keys to split.

Return type `jax.Array`

Returns An array of shape `[num, 2]` unique (within a transformed function) JAX rng keys that can be used with APIs such as `jax.random.uniform()`.

maybe_next_rng_key

`haiku.maybe_next_rng_key()`

`next_rng_key()` if random numbers are available, else `None`.

Return type `Optional[PRNGKey]`

reserve_rng_keys

`haiku.reserve_rng_keys(num)`

Pre-allocate some number of JAX RNG keys.

See `next_rng_key()`.

This API offers a way to micro-optimize how RNG keys are split when using Haiku. It is unlikely that you need it unless you find compilation time of your `init` function to be a problem, or you sample a lot of random numbers in apply.

```
>>> hk.reserve_rng_keys(2) # Pre-allocate 2 keys for us to consume.
>>> _ = hk.next_rng_key()  # Takes the first pre-allocated key.
>>> _ = hk.next_rng_key()  # Takes the second pre-allocated key.
>>> _ = hk.next_rng_key()  # Splits a new key.
```

Parameters `num (int)` – The number of JAX rng keys to allocate.

with_rng

`haiku.with_rng(key)`

Provides a new sequence for `next_rng_key()` to draw from.

When `next_rng_key()` is called, it draws a new key from the `PRNGSequence` defined by the input key to the transformed function. This context manager overrides the sequence for the duration of the scope.

```
>>> with hk.with_rng(jax.random.PRNGKey(428)):
...     s = jax.random.uniform(hk.next_rng_key(), ())
>>> print("{:.1f}".format(s))
0.5
```

Parameters `key (PRNGKey)` – The key to seed the sequence with.

Returns Context manager under which the given sequence is active.

maybe_get_rng_sequence_state

`haiku.maybe_get_rng_sequence_state()`

Returns the internal state of the PRNG sequence.

Return type `Optional[PRNGSequenceState]`

Returns The internal state if random numbers are available, else `None`.

replace_rng_sequence_state

`haiku.replace_rng_sequence_state(state)`

Replaces the internal state of the PRNG sequence with the given state.

Parameters `state` (*PRNGSequenceState*) – The new internal state or `None`.

Raises `MissingRNGError` – If random numbers aren't available.

Type Hints

<i>LSTMState</i> (hidden, cell)	An LSTM core state consists of hidden and cell vectors.
<i>Params</i>	A Mapping is a generic container for associating key/value pairs.
<i>MutableParams</i>	A MutableMapping is a generic container for associating key/value pairs.
<i>State</i>	A Mapping is a generic container for associating key/value pairs.
<i>MutableState</i>	A MutableMapping is a generic container for associating key/value pairs.
<i>Transformed</i> (init, apply)	Holds a pair of pure functions.
<i>TransformedWithState</i> (init, apply)	Holds a pair of pure functions.
<i>MultiTransformed</i> (init, apply)	Holds a collection of pure functions.
<i>MultiTransformedWithState</i> (init, apply)	Holds a collection of pure functions.
<i>ModuleProtocol</i> (*args, **kwargs)	Protocol for Module like types.
<i>SupportsCall</i> (*args, **kwargs)	Protocol for Module like types that are Callable.

LSTMState

class `haiku.LSTMState`(*hidden: jax.Array, cell: jax.Array*)

An LSTM core state consists of hidden and cell vectors.

hidden

Hidden state.

Type `jax.Array`

cell

Cell state.

Type `jax.Array`

Params

`haiku.Params`

alias of `collections.abc.Mapping[str, collections.abc.Mapping[str, jax.Array]]`

MutableParams

haiku.MutableParams

alias of collections.abc.MutableMapping[str, collections.abc.MutableMapping[str, jax.Array]]

State

haiku.State

alias of collections.abc.Mapping[str, collections.abc.Mapping[str, jax.Array]]

MutableState

haiku.MutableState

alias of collections.abc.MutableMapping[str, collections.abc.MutableMapping[str, jax.Array]]

Transformed

class haiku.Transformed(*init: Callable[..., hk.MutableParams], apply: Callable[..., Any]*)

Holds a pair of pure functions.

init

A pure function: `params = init(rng, *a, **k)`

Type Callable[..., hk.MutableParams]

apply

A pure function: `out = apply(params, rng, *a, **k)`

Type Callable[..., Any]

TransformedWithState

class haiku.TransformedWithState(*init: Callable[..., tuple[hk.MutableParams, hk.MutableState]], apply: Callable[..., tuple[Any, hk.MutableState]]*)

Holds a pair of pure functions.

init

A pure function: `params, state = init(rng, *a, **k)`

Type Callable[..., tuple[hk.MutableParams, hk.MutableState]]

apply

A pure function: `out, state = apply(params, state, rng, *a, **k)`

Type Callable[..., tuple[Any, hk.MutableState]]

MultiTransformed

class haiku.**MultiTransformed**(*init: Callable[...], hk.MutableParams], apply: Any*)

Holds a collection of pure functions.

init

A pure function: `params = init(rng, *a, **k)`

Type Callable[...], hk.MutableParams]

apply

A JAX tree of pure functions each with the signature: `out = apply(params, rng, *a, **k)`.

Type Any

See also:

- *Transformed*: Single apply variant of multi-transform.
- *MultiTransformedWithState*: Multi apply with state variant.

MultiTransformedWithState

class haiku.**MultiTransformedWithState**(*init: Callable[...], tuple[hk.MutableParams, hk.MutableState]], apply: Any*)

Holds a collection of pure functions.

init

A pure function: `params, state = init(rng, *a, **k)`

Type Callable[...], tuple[hk.MutableParams, hk.MutableState]]

apply

A JAX tree of pure functions each with the signature: `out, state = apply(params, state, rng, *a, **k)`.

Type Any

See also:

- *TransformedWithState*: Single apply variant of multi-transform.
- *MultiTransformed*: Multi apply with state variant.

ModuleProtocol

class haiku.**ModuleProtocol**(*args, **kwargs)

Protocol for Module like types.

SupportsCall

```
class haiku.SupportsCall(*args, **kwargs)
```

Protocol for Module like types that are Callable.

Being a protocol means you don't need to explicitly extend this type in order to support instance checks with it. For example, `Linear` only extends `Module`, however since it conforms (e.g. implements `__call__`) to this protocol you can instance check using it:

```
>>> assert isinstance(hk.Linear(1), hk.SupportsCall)
```

1.3.2 Flax Interop

Haiku inside Flax

Module

```
class haiku.experimental.flax.Module(transformed, parent=<flax.linen.module._Sentinel object>,
                                     name=None)
```

A Flax `nn.Module` that runs a Haiku transformed function.

This type is designed to make it easy to take a Haiku transformed function and/or a Haiku module and use it inside a program that otherwise uses Flax.

Given a Haiku transformed function

```
>>> def f(x):
...     return hk.Linear(1)(x)
>>> f = hk.transform(f)
```

You can convert it into a Flax module using:

```
>>> mod = hk.experimental.flax.Module(f)
```

Calling this module is the same as calling any regular Flax module:

```
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 1])
>>> variables = mod.init(rng, x)
>>> out = mod.apply(variables, x)
```

If you just want to convert a Haiku module class such that it can be used with Flax you can use the `create` class method:

```
>>> mod = hk.experimental.flax.Module.create(hk.Linear, 1)
>>> variables = mod.init(rng, x)
>>> out = mod.apply(variables, x)
```

flatten_flax_to_haiku

`haiku.experimental.flax.flatten_flax_to_haiku(collection)`

Flattens a Flax variable collection (e.g. params) to a Haiku dict.

Return type HaikuParamsOrState

Flax inside Haiku

lift

`haiku.experimental.flax.lift(mod, *, name)`

Lifts a flax `nn.Module` into a Haiku transformed function.

For a Flax Module (e.g. `mod = nn.Dense(10)`), `mod = lift(mod)` allows you to run the call method of the module as if the module was a regular Haiku module.

Parameters and state from the Flax module are registered with Haiku and become part of the params/state dictionaries (as returned from `init/apply`).

```

>>> def f(x):
...     # Create and "lift" a Flax module.
...     mod = hk.experimental.flax.lift(nn.Dense(300), name='dense')
...     x = mod(x)                # Any params/state will be registered
...                               # with Haiku when applying the module.
...     x = jax.nn.relu(x)
...     x = hk.nets.MLP([100, 10]) # You can of course mix Haiku modules in.
...     return x
>>> f = hk.transform(f)
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 1])
>>> params = f.init(rng, x)
>>> out = f.apply(params, None, x)

```

Parameters

- **mod** (*nn.Module*) – Any Flax `nn.Module` instance.
- **name** (*str*) – Name scope to prefix entries in the outer params/state dict.

Return type Callable[*...*, Any]

Returns A function that when applied calls the call method of the given Flax module and returns its output. As a side effect of calling the module any module parameters and state variables are registered with Haiku.

1.3.3 Advanced State Management

Lifting

<code>lift(init_fn, *[, allow_reuse, name])</code>	Registers parameters from an inner init function in an outer transform.
<code>lift_with_state(init_fn, *[, allow_reuse, name])</code>	Registers params and state from an init function in an outer transform.
<code>transparent_lift(init_fn, *[, allow_reuse])</code>	Registers parameters in an outer transform without adding a name scope.
<code>transparent_lift_with_state(init_fn, *[, ...])</code>	Registers params and state in an outer transform without adding scope.
<code>LiftWithStateUpdater(name)</code>	Handles updating the state for a <code>lift_with_state</code> computation.

lift

`haiku.lift`(*init_fn*, *, *allow_reuse=False*, *name='lifted'*)

Registers parameters from an inner init function in an outer transform.

HINT: `lift()` is for when you want to make non-trivial use of JAX transforms (e.g. `jax.vmap`) **inside** of a `transform()` or `transform_with_state()`. We generally recommend trying to use JAX transforms on the pure functions returned by `transform()`, in which case you do not need `lift()`.

Use `lift()` when nesting Haiku transforms to register the parameters of the inner transform in any outer transform. This is mainly useful when using JAX functions inside of a Haiku module (eg. using `jax.vmap` on a layer). See <https://dm-haiku.readthedocs.io/en/latest/notebooks/transforms.html#Using-hk.lift> for more explanation of when to use `lift()`. (If you're not using JAX functions inside of a module or don't need access to your parameters inside of a transform, you probably don't need to use `lift()`)

Must be called inside `transform()`, and be passed the `init` member of a `Transformed`.

During init, the returned callable will run the given `init_fn`, and include the resulting params in the outer transform's dictionaries. During apply, the returned callable will instead pull the relevant parameters from the outer transform's dictionaries.

By default, users must ensure that the given `init` does not accidentally catch modules from an outer `transform()` via functional closure. If this behavior is desirable, set `allow_reuse` to `True`.

Example:

A common usage of `lift()` is to use JAX transformations like `vmap` in non-trivial ways, inside a `transform()`. For example, we can use `lift()` and `jax.vmap` to create an ensemble.

First we'll create a helper function that uses `lift()` to apply `vmap` to our model. As you can see from the comments, we are using `vmap` to change how parameters should be created (in this case we create a unique set of parameters for each member of the ensemble) and we change how apply works (we "map" the parameters meaning JAX will compute the forward pass separately, in parallel, for each member of the ensemble):

```
>>> def create_ensemble(model, size: int):
...     init_rng = hk.next_rng_keys(size) if hk.running_init() else None
...     model = hk.transform(model)
...     # in_axes: rng is mapped, data is not.
...     init_model = jax.vmap(model.init, in_axes=(0, None))
...     # Use hk.lift to "lift" parameters created by `init_model` into the
```

(continues on next page)

(continued from previous page)

```

... # outer transform.
... init_model = hk.lift(init_model, name="ensemble")
... def ensemble(x):
...     params = init_model(init_rng, x)
...     # in_axes: params are mapped, rng/data are not.
...     return jax.vmap(model.apply, in_axes=(0, None, None))(params, None, x)
...     return ensemble

```

We can now use this function to ensemble any Haiku module(s), inside of a transform. First we define a function for each member of the ensemble:

```

>>> def member_fn(x):
...     return hk.nets.MLP([300, 100, 10])(x)

```

Secondly we can combine our two functions, inside a `transform()` to create an ensemble:

```

>>> def f(x):
...     ensemble = create_ensemble(member_fn, size=4)
...     x = ensemble(x)
...     # You could create other modules here which were not ensembled.
...     return x
>>> f = hk.transform(f)

```

When we initialize the network, our ensemble member's parameters have a leading dimension the size of the ensemble:

```

>>> rng = jax.random.PRNGKey(777)
>>> x = jnp.ones([32, 128])
>>> params = f.init(rng, x)
>>> jax.tree_util.tree_map(lambda x: x.shape, params)
{'ensemble/mlp/~linear_0': {'b': (4, 300), 'w': (4, 128, 300)},
 'ensemble/mlp/~linear_1': {'b': (4, 100), 'w': (4, 300, 100)},
 'ensemble/mlp/~linear_2': {'b': (4, 10), 'w': (4, 100, 10)}}

```

When we apply the network, we get an output for each member of the ensemble for the entire batch:

```

>>> y = f.apply(params, None, x)
>>> y.shape
(4, 32, 10)

```

Parameters

- **init_fn** (*Callable*[..., *hk.Params*]) – The init function from an *Transformed*.
- **allow_reuse** (*bool*) – Allows lifted parameters and state to be reused from the outer *transform()*. This can be desirable when using *lift* within control flow (e.g. *hk.scan*).
- **name** (*str*) – A string name to prefix parameters with.

Return type *Callable*[..., *hk.Params*]

Returns A callable that during *init* injects parameter values into the outer context and during *apply* retrieves parameters from the outer context. In both cases returns parameter values to be used with an *apply* function.

See also:

- `lift_with_state()`: Register params and state with an outer transform.
- `transparent_lift()`: Register params with an outer transform without a namespace.
- `transparent_lift_with_state()`: Register params and state with an outer transform without a namespace.

lift_with_state

`haiku.lift_with_state(init_fn, *, allow_reuse=False, name='lifted')`

Registers params and state from an init function in an outer transform.

See `lift()` for more context on when to use `lift`.

This function returns two objects. The first is a callable that runs your init function with slightly different behaviour based on if it's run during init vs. apply time. The second is an updater that can be used to pass updated state values that result from running your apply function. See later in the docs for a worked example.

During init, the returned callable will run the given `init_fn`, and include the resulting params/state in the outer transform's dictionaries. During apply, the returned callable will instead pull the relevant params/state from the outer transform's dictionaries.

Must be called inside `transform_with_state()`, and be passed the `init` member of a `TransformedWithState`.

By default, users must ensure that the given `init` does not accidentally catch modules from an outer `transform_with_state()` via functional closure. If this behavior is desirable, set `allow_reuse` to `True`.

Example

```
>>> def g(x):
...     return hk.nets.ResNet50(1)(x, True)
>>> g = hk.transform_with_state(g)
>>> params_and_state_fn, updater = (
...     hk.lift_with_state(g.init, name='f_lift'))
>>> init_rng = hk.next_rng_key() if hk.running_init() else None
>>> x = jnp.ones([1, 224, 224, 3])
>>> params, state = params_and_state_fn(init_rng, x)
>>> out, state = g.apply(params, state, None, x)
>>> updater.update(state)
```

Parameters

- `init_fn` (`Callable[... , tuple[hk.Params, hk.State]]`) – The `init` function from an `TransformedWithState`.
- `allow_reuse` (`bool`) – Allows lifted parameters and state to be reused from the outer `transform_with_state()`. This can be desirable when using `lift_with_state` within control flow (e.g. `hk.scan`).
- `name` (`str`) – A string name to prefix parameters with.

Return type `tuple[Callable[... , tuple[hk.Params, hk.State]], LiftWithStateUpdater]`

Returns A callable that during `init` injects parameter values into the outer context and during `apply` reuses parameters from the outer context. In both cases returns parameter values to be used with

an `apply` function. The `init` function additionally returns an object used to update the outer context with new state after `apply` is called.

See also:

- `lift()`: Register parameters with an outer transform.
- `transparent_lift()`: Register parameters with an outer transform without a namespace.
- `transparent_lift_with_state()`: Register parameters and state with an outer transform without a namespace.

transparent_lift

`haiku.transparent_lift`(*init_fn*, *, *allow_reuse=False*)

Registers parameters in an outer transform without adding a name scope.

Functionally this is equivalent to `lift()` but without automatically adding an additional variable scoping. Note that closing over a module from an outer scope is disallowed.

See `lift()` for more context on when to use `lift`.

Parameters

- `init_fn` (*Callable*[..., *hk.Params*]) – The `init` function from an *Transformed*.
- `allow_reuse` (*bool*) – Allows lifted parameters to be reused from the outer `transform_with_state()`. This can be desirable when e.g. within control flow (e.g. `hk.scan`).

Return type *Callable*[..., *hk.Params*]

Returns A callable that during `init` injects parameter values into the outer context and during `apply` reuses parameters from the outer context. In both cases returns parameter values to be used with an `apply` function.

See also:

- `lift()`: Register params with an outer transform.
- `lift_with_state()`: Register params and state with an outer transform.
- `transparent_lift_with_state()`: Register params and state with an outer transform without a namespace.

transparent_lift_with_state

`haiku.transparent_lift_with_state`(*init_fn*, *, *allow_reuse=False*)

Registers params and state in an outer transform without adding scope.

Functionally this is equivalent to `lift_with_state()` but without automatically adding an additional variable scoping.

See `lift_with_state()` for more context on when to use `lift_with_state`.

Parameters

- `init_fn` (*Callable*[..., *tuple*[*hk.Params*, *hk.State*]]) – The `init` function from an *TransformedWithState*.

- **allow_reuse** (*bool*) – Allows lifted parameters and state to be reused from the outer `transform_with_state()`. This can be desirable when e.g. within control flow (e.g. `hk.scan`).

Return type `tuple[Callable[...], tuple[hk.Params, hk.State]], LiftWithStateUpdater]`

Returns A callable that during `init` injects parameter values into the outer context and during `apply` reuses parameters from the outer context. In both cases returns parameter values to be used with an `apply` function. The `init` function additionally returns an object used to update the outer context with new state after `apply` is called.

See also:

- `lift()`: Register params with an outer transform.
- `lift_with_state()`: Register params and state with an outer transform.
- `transparent_lift()`: Register params with an outer transform without a namespace.

LiftWithStateUpdater

class `haiku.LiftWithStateUpdater(name)`

Handles updating the state for a `lift_with_state` computation.

Layer Stack

<code>layer_stack(num_layers[, ...])</code>	Utility to wrap a Haiku function and recursively apply it to an input.
<code>LayerStackTransparencyMapping(*args, **kwargs)</code>	Module name mapping for transparent <code>layer_stack</code> .

layer_stack

class `haiku.layer_stack(num_layers, with_per_layer_inputs=False, unroll=1, pass_reverse_to_layer_fn=False, transparent=False, transparency_map=None, name=None)`

Utility to wrap a Haiku function and recursively apply it to an input.

This can be used to improve model compile times.

A function is valid if it uses only explicit position parameters, and its return type matches its input type. The position parameters can be arbitrarily nested structures with `jax.Array` at the leaf nodes. Note that kwargs are not supported, neither are functions with variable number of parameters (specified by `*args`).

Note that `layer_stack` cannot at the moment be used with functions that build Haiku modules with state.

If `with_per_layer_inputs=False` then the new, wrapped function can be understood as performing the following:

```
>>> f = lambda x: x+1
>>> num_layers = 4
>>> x = 0
>>> for i in range(num_layers):
```

(continues on next page)

(continued from previous page)

```
... x = f(x)
>>> x
4
```

And if `with_per_layer_inputs=True`, assuming `f` takes two arguments on top of `x`:

```
>>> f = lambda x, y0, y1: (x+1, y0+y1)
>>> num_layers = 4
>>> x = 0
>>> ys_0 = [1, 2, 3, 4]
>>> ys_1 = [5, 6, 7, 8]
>>> zs = []
>>> for i in range(num_layers):
...     x, z = f(x, ys_0[i], ys_1[i])
...     zs.append(z)
>>> x, zs
(4, [6, 8, 10, 12])
```

The code using `layer_stack` for the above function would be:

```
>>> f = lambda x, y0, y1: (x+1, y0+y1)
>>> num_layers = 4
>>> x = 0
>>> ys_0 = jnp.array([1, 2, 3, 4])
>>> ys_1 = jnp.array([5, 6, 7, 8])
>>> stack = hk.layer_stack(num_layers, with_per_layer_inputs=True)
>>> x, zs = stack(f)(x, ys_0, ys_1)
>>> print(x, zs)
4 [ 6  8 10 12]
```

Check the tests in `layer_stack_test.py` for further examples.

Crucially, any parameters created inside `f` will not be shared across iterations.

Parameters

- **num_layers** (*int*) – The number of times to iterate the wrapped function.
- **with_per_layer_inputs** – Whether or not to pass per-layer inputs to the wrapped function.
- **unroll** (*int*) – the unroll used by scan.
- **pass_reverse_to_layer_fn** (*bool*) – Whether or not to pass the `reverse` keyword to the function `f`, so that it is aware if the layer stack is being run forward or in reverse (and the underlying scan). To run the layer stack in reverse you need to pass in `reverse=True` to the call to the layer stack.
- **transparent** (*bool*) – Whether to apply `layer_stack` transparently. When this is `True`, and a correct `transparency_map` is provided, the parameters are generated in such a way that `layer_stack` can be replaced by a regular for loop without changing the parameter tree.
- **transparency_map** (*Optional[LayerStackTransparencyMapping]*) – How to map stacked module names to flat names and reverse. See `LayerStackTransparencyMapping` and `layer_stack_test.py` for an example.
- **name** (*Optional[str]*) – name of the Haiku context.

Returns Callable that will produce a layer stack when called with a valid function.

LayerStackTransparencyMapping

class haiku.LayerStackTransparencyMapping(*args, **kwargs)

Module name mapping for transparent layer_stack.

Naming

<code>name_scope(name, *, [method_name])</code>	Context manager which adds a prefix to all new modules, params or state.
<code>current_name()</code>	Returns the currently active module name.
<code>DO_NOT_STORE</code>	Causes a parameter or state value to not be stored.
<code>get_params()</code>	Returns the parameters for the current <code>transform()</code> .
<code>get_current_state()</code>	Returns the current state for the current <code>transform_with_state()</code> .
<code>get_initial_state()</code>	Returns the initial state for the current <code>transform_with_state()</code> .
<code>force_name(name)</code>	Forces Haiku to use this name, ignoring all context information.
<code>name_like(method_name)</code>	Allows a method to be named like some other method.
<code>transparent(method)</code>	Decorator to wrap a method, preventing automatic variable scope wrapping.

name_scope

`haiku.name_scope(name, *, method_name='__call__')`

Context manager which adds a prefix to all new modules, params or state.

```
>>> with hk.name_scope("my_name_scope"):
...     net = hk.Linear(1, name="my_linear")
>>> net.module_name
'my_name_scope/my_linear'
```

When used inside a module, any submodules, parameters or state created inside the name scope will have a prefix added to their names:

```
>>> class MyModule(hk.Module):
...     def __call__(self, x):
...         with hk.name_scope("my_name_scope"):
...             submodule = hk.Linear(1, name="submodule")
...             w = hk.get_parameter("w", [], init=jnp.ones)
...         return submodule(x) + w
```

```
>>> f = hk.transform(lambda x: MyModule()(x))
>>> params = f.init(jax.random.PRNGKey(42), jnp.ones([1, 1]))
>>> jax.tree_util.tree_map(jnp.shape, params)
{'my_module/my_name_scope': {'w': ()},
 'my_module/my_name_scope/submodule': {'b': (1,), 'w': (1, 1)}}
```

Name scopes are very similar to putting all of the code inside the context manager inside a method on a *Module* with the name you provide. Behind the scenes this is precisely how name scopes are implemented.

If you are familiar with TensorFlow then Haiku's `name_scope()` is similar to `tf.variable_scope(..)` in TensorFlow 1 and `tf.name_scope(..)` in TensorFlow 1 and 2 in that it changes the names associated with modules, parameters and state.

Parameters

- **name** (*str*) – The name scope to use (e.g. "foo" or "foo/bar").
- **method_name** (*str*) – (Advanced uses only). Since name scopes are equivalent to calling methods on modules the method name attribute allows you to specify which method name you want to simulate. Most users should leave this as the default value ("`__call__`").

Return type ContextManager[None]

Returns A single use context manager that when active prefixes new modules, parameters or state with the given name.

current_name

`haiku.current_name()`

Returns the currently active module name.

Outside of a Haiku module (but inside a Haiku transform) this will return `~` which matches the key in the `params/state` dict where top level values are stored.

```
>>> hk.current_name()
'~'
```

Inside a module this returns the current module name:

```
>>> class ExampleModule(hk.Module):
...     def __call__(self):
...         return hk.current_name()
>>> ExampleModule()()
'example_module'
```

Inside a name scope this returns the current name scope:

```
>>> with hk.name_scope('example_name_scope'):
...     print(hk.current_name())
example_name_scope
```

Return type `str`

Returns The currently active module or name scope name. If modules or name scopes are in use returns `~`.

DO_NOT_STORE

`haiku.DO_NOT_STORE` = <haiku._src.base.DoNotStore object>

Causes a parameter or state value to not be stored.

By default, Haiku will put the value returned from `get_parameter()`, `get_state()` and `set_state()` into the dictionaries returned by `init`. This is not always desirable.

For example, a user may want to have part of their network come from a pretrained checkpoint, and they may want to freeze those values (aka. have them not appear in the params dict passed later to `grad`). You can achieve this by manipulating the params dict, however sometimes it is more convenient to do this using custom creators/getters/setters.

Consider the following function:

```
>>> def f(x):
...     x = hk.Linear(300, name='torso')(x)
...     x = hk.Linear(10, name='tail')(x)
...     return x
```

Imagine you have a pre-trained set of weights for the torso:

```
>>> pretrained = {'torso': {'w': jnp.ones([28 * 28, 300]),
...                             'b': jnp.ones([300])}}
```

First we define a creator, that tells Haiku to not store any parameters that are part of the pretrained dict:

```
>>> def my_creator(next_creator, shape, dtype, init, context):
...     if context.module_name in pretrained:
...         return hk.DO_NOT_STORE
...     return next_creator(shape, dtype, init)
```

Then we need a getter that provides the parameter value from the pretrained dict:

```
>>> def my_getter(next_getter, value, context):
...     if context.module_name in pretrained:
...         assert value is hk.DO_NOT_STORE
...         value = pretrained[context.module_name][context.name]
...     return next_getter(value)
```

Finally we'll wrap our function in context managers activating our creator and getter:

```
>>> def f_with_pretrained_torso(x):
...     with hk.custom_creator(my_creator), \
...         hk.custom_getter(my_getter):
...         return f(x)
```

You can see that when we run our function we only get parameters from modules that were not in the pretrained dict:

```
>>> f_with_pretrained_torso = hk.transform(f_with_pretrained_torso)
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 28 * 28])
>>> params = f_with_pretrained_torso.init(rng, x)
>>> assert list(params) == ['tail']
```

This value can be used in initialisers, `custom_creator()` or `custom_setter()`.

get_params

haiku.get_params()

Returns the parameters for the current `transform()`.

```
>>> def report(when):
...     shapes = jax.tree_util.tree_map(jnp.shape, hk.get_params())
...     print(f'{when}: {shapes}')
>>> def f(x):
...     report('Before call')
...     x = hk.Linear(1)(x)
...     report('After call')
...     return x
>>> f = hk.transform(f)
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([1, 1])
```

During `init` the parameters dictionary will get populated as modules are called:

```
>>> params = f.init(rng, x)
Before call: {}
After call: {'linear': {'b': (1,), 'w': (1, 1)}}
```

During `apply` the parameters dictionary will remain unchanged:

```
>>> _ = f.apply(params, None, x)
Before call: {'linear': {'b': (1,), 'w': (1, 1)}}
```

NOTE: Does not run `custom_getters()` or parameter initializers.

Return type Params

Returns A copy of the parameters dictionary. During `init` this dictionary will be populated with any parameters that have been created so far. During `apply` this will contain all parameters of all modules (the params dict does not change during `apply`).

See also:

- `get_initial_state()`: The initial state for the function.
- `get_current_state()`: The current state for the function.

get_current_state

`haiku.get_current_state()`

Returns the current state for the current `transform_with_state()`.

Example:

```

>>> def report(when):
...     state = jax.tree_util.tree_map(int, hk.get_current_state())
...     print(f'{when}: {state}')
>>> def f():
...     report('Before get_state')
...     x = hk.get_state('x', [], init=jnp.zeros)
...     report('After get_state')
...     hk.set_state('x', x + 1)
...     report('After set_state')
>>> f = hk.transform_with_state(f)

```

During `init`, the most recently set value (either directly via `set_state()` or via the `init` argument to `get_state()`) will be returned:

```

>>> _, state = f.init(None)
Before get_state: {}
After get_state: {'~': {'x': 0}}
After set_state: {'~': {'x': 1}}

```

During `apply` the most recently set value will be used, if no value has been set then the value that is passed into `apply` will be used:

```

>>> state = {'~': {'x': 10}}
>>> _ = f.apply({}, state, None)
Before get_state: {'~': {'x': 10}}
After get_state: {'~': {'x': 10}}
After set_state: {'~': {'x': 11}}

```

NOTE: Does not run `custom_getters()` or state initializers.

Return type State

Returns A copy of the state dictionary that would be returned from `init` or `apply`.

See also:

- `get_params()`: The current parameters for the function.
- `get_initial_state()`: The initial state for the function.

get_initial_state

haiku.get_initial_state()

Returns the initial state for the current `transform_with_state()`.

Example:

```

>>> def report(when):
...     state = jax.tree_util.tree_map(int, hk.get_initial_state())
...     print(f'{when}: {state}')
>>> def f():
...     report('Before get_state')
...     x = hk.get_state('x', [], init=jnp.zeros)
...     report('After get_state')
...     hk.set_state('x', x + 1)
...     report('After set_state')
>>> f = hk.transform_with_state(f)

```

During init, the first set value (either directly via `set_state()` or via the `init` argument to `get_state()`) will be returned:

```

>>> _, state = f.init(None)
Before get_state: {}
After get_state: {'~': {'x': 0}}
After set_state: {'~': {'x': 0}}

```

During apply the value passed into the apply function will be used:

```

>>> state = {'~': {'x': 10}}
>>> _ = f.apply({}, state, None)
Before get_state: {'~': {'x': 10}}
After get_state: {'~': {'x': 10}}
After set_state: {'~': {'x': 10}}

```

NOTE: Does not run `custom_getters()` or state initializers.

Return type State

Returns A copy of the state dictionary that would be returned from `init` or passed into `apply`.

See also:

- `get_params()`: The current parameters for the function.
- `get_current_state()`: The current state for the function.

force_name

haiku.force_name(name)

Forces Haiku to use this name, ignoring all context information.

NOTE: This method is intended for advanced use cases only and should be avoided whenever possible as it effectively enforces a singleton pattern when setting absolute names.

Haiku names modules according to where they are created (e.g. the stack of modules that created them, or the current `name_scope()`). This function allows you to create modules that ignore all of this and have precisely the name you provide.

This might be useful in the case that you have two modules and you want to force them to share parameters:

```
>>> mod0 = hk.Linear(1)
>>> some_hyperparameter = True
>>> if some_hyperparameter:
...     # Force mod1 and mod0 to have shared weights.
...     mod1 = hk.Linear(1, name=hk.force_name(mod0.module_name))
... else:
...     # mod0 and mod1 are independent.
...     mod1 = hk.Linear(1)
```

(A simpler version of this snippet would do `mod1 = mod0` instead of using `force_name`, however in real examples it can be simpler to use `force_name`, especially in cases where you may not have access to the module instance without lots of plumbing, but getting the module name is easy [e.g. it is a hyperparameter]).

Parameters `name` (*str*) – String name for the module. For example "foo" or "foo/bar".

Return type *str*

Returns A value suitable to pass into the `name` argument of any Haiku module constructor.

name_like

`haiku.name_like`(*method_name*)

Allows a method to be named like some other method.

In Haiku submodules are named based on the name of their parent module and the method in which they are created. When refactoring code it may be desirable to maintain previous names in order to keep checkpoint compatibility, this can be achieved using `name_like()`.

As an example, consider the following toy autoencoder:

```
>>> class Autoencoder(hk.Module):
...     def __call__(self, x):
...         z = hk.Linear(10, name="enc")(x) # name: autoencoder/enc
...         y = hk.Linear(10, name="dec")(z) # name: autoencoder/dec
...         return y
```

If we want to refactor this such that users can encode or decode, we would create two methods (`encode`, `decode`) which would create and apply our modules. In order to retain checkpoint compatibility with the original module we can use `name_like()` to name those submodules as if they were created inside `__call__`:

```
>>> class Autoencoder(hk.Module):
...     @hk.name_like("__call__")
...     def encode(self, x):
...         return hk.Linear(10, name="enc")(x) # name: autoencoder/enc
...
...     @hk.name_like("__call__")
...     def decode(self, z):
...         return hk.Linear(10, name="dec")(z) # name: autoencoder/dec
...
...     def __call__(self, x):
...         return self.decode(self.encode(x))
```

One sharp edge is if users rely on Haiku's numbering to take care of giving unique names and refactor using `name_like()`. For example when refactoring the following:

```
>>> class Autoencoder(hk.Module):
...     def __call__(self, x):
...         y = hk.Linear(10)(z) # name: autoencoder/linear_1
...         z = hk.Linear(10)(x) # name: autoencoder/linear
...         return y
```

To use `name_like()`, the unnamed linear modules in encode/decode will end up with the same name (both: autoencoder/linear) because module numbering is only applied within a method:

```
>>> class Autoencoder(hk.Module):
...     @hk.name_like("__call__")
...     def encode(self, x):
...         return hk.Linear(10)(x) # name: autoencoder/linear
...
...     @hk.name_like("__call__")
...     def decode(self, z):
...         return hk.Linear(10)(z) # name: autoencoder/linear <-- NOT INTENDED
```

To fix this case you need to explicitly name the modules within the method with their former name:

```
>>> class Autoencoder(hk.Module):
...     @hk.name_like("__call__")
...     def encode(self, x):
...         return hk.Linear(10, name="linear")(x) # name: autoencoder/linear
...
...     @hk.name_like("__call__")
...     def decode(self, z):
...         return hk.Linear(10, name="linear_1")(z) # name: autoencoder/linear_1
```

Parameters `method_name` (*str*) – The name of a method whose name we should adopt. This method does not actually have to be defined on the class.

Return type Callable[[T], T]

Returns A decorator that when applied to a method marks it as having a different name.

transparent

`haiku.transparent` (*method*)

Decorator to wrap a method, preventing automatic variable scope wrapping.

By default, all variables and modules created in a method are scoped by the module and method names. This is undesirable in some cases. Any method decorated with `transparent()` will create variables and modules in the scope in which it was called.

Parameters `method` (*T*) – the method to wrap.

Return type T

Returns The method, with a flag indicating no name scope wrapping should occur.

Visualisation

<code>to_dot(fun)</code>	Converts a function using Haiku modules to a dot graph.
--------------------------	---

to_dot

`haiku.to_dot(fun)`

Converts a function using Haiku modules to a dot graph.

To view the resulting graph in Google Colab or an iPython notebook use the `graphviz` package:

```
dot = hk.to_dot(f)(x)
import graphviz
graphviz.Source(dot)
```

Parameters `fun` (*Callable*[... , *Any*]) – A function using Haiku modules.

Return type *Callable*[... , str]

Returns A function that returns the source code string to a graphviz graph describing the operations executed by the given function clustered by Haiku module.

See also:

`abstract_to_dot()`: Generates a graphviz graph using abstract inputs.

1.3.4 Common Modules

Linear

<code>Linear(output_size[, with_bias, w_init, ...])</code>	Linear module.
<code>Bias([output_size, bias_dims, b_init, name])</code>	Adds a bias to inputs.

Linear

class `haiku.Linear(output_size, with_bias=True, w_init=None, b_init=None, name=None)`

Linear module.

`__init__` (*output_size*, *with_bias=True*, *w_init=None*, *b_init=None*, *name=None*)

Constructs the Linear module.

Parameters

- **output_size** (*int*) – Output dimensionality.
- **with_bias** (*bool*) – Whether to add a bias to the output.
- **w_init** (*Optional*[*hk.initializers.Initializer*]) – Optional initializer for weights. By default, uses random values from truncated normal, with stddev $1 / \sqrt{\text{fan_in}}$. See <https://arxiv.org/abs/1502.03167v3>.
- **b_init** (*Optional*[*hk.initializers.Initializer*]) – Optional initializer for bias. By default, zero.

- **name** (*Optional[str]*) – Name of the module.

`__call__` (*inputs, *, precision=None*)

Computes a linear transform of the input.

Return type `jax.Array`

Bias

class `haiku.Bias` (*output_size=None, bias_dims=None, b_init=None, name=None*)

Adds a bias to inputs.

Example Usage:

```
>>> N, H, W, C = 1, 2, 3, 4
>>> x = jnp.ones([N, H, W, C])
>>> scalar_bias = hk.Bias(bias_dims=[])
>>> scalar_bias_output = scalar_bias(x)
>>> assert scalar_bias.bias_shape == ()
```

Create a bias over all non-minibatch dimensions:

```
>>> all_bias = hk.Bias()
>>> all_bias_output = all_bias(x)
>>> assert all_bias.bias_shape == (H, W, C)
```

Create a bias over the last non-minibatch dimension:

```
>>> last_bias = hk.Bias(bias_dims=[-1])
>>> last_bias_output = last_bias(x)
>>> assert last_bias.bias_shape == (C,)
```

Create a bias over the first non-minibatch dimension:

```
>>> first_bias = hk.Bias(bias_dims=[1])
>>> first_bias_output = first_bias(x)
>>> assert first_bias.bias_shape == (H, 1, 1)
```

Subtract and later add the same learned bias:

```
>>> bias = hk.Bias()
>>> h1 = bias(x, multiplier=-1)
>>> h2 = bias(x)
>>> h3 = bias(x, multiplier=-1)
>>> reconstructed_x = bias(h3)
>>> assert (x == reconstructed_x).all()
```

`__init__` (*output_size=None, bias_dims=None, b_init=None, name=None*)

Constructs a Bias module that supports broadcasting.

Parameters

- **output_size** (*Optional[Sequence[int]]*) – Output size (output shape without batch dimension). If `output_size` is left as `None`, the size will be directly inferred by the input.

- **bias_dims** (*Optional[Sequence[int]]*) – Sequence of which dimensions to retain from the input shape when constructing the bias. The remaining dimensions will be broadcast over (given size of 1), and leading dimensions will be removed completely. See class doc for examples.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for the bias. Default to zeros.
- **name** (*Optional[str]*) – Name of the module.

`__call__(inputs, multiplier=None)`

Adds bias to `inputs` and optionally multiplies by `multiplier`.

Parameters

- **inputs** (*jax.Array*) – A Tensor of size `[batch_size, input_size1, ...]`.
- **multiplier** (*Optional[Union[float, jax.Array]]*) – A scalar or Tensor which the bias term is multiplied by before adding it to `inputs`. Anything which works in the expression `bias * multiplier` is acceptable here. This may be useful if you want to add a bias in one place and subtract the same bias in another place via `multiplier=-1`.

Return type `jax.Array`

Returns A Tensor of size `[batch_size, input_size1, ...]`.

Pooling

<code>avg_pool(value, window_shape, strides, padding)</code>	Average pool.
<code>AvgPool(window_shape, strides, padding[, ...])</code>	Average pool.
<code>max_pool(value, window_shape, strides, padding)</code>	Max pool.
<code>MaxPool(window_shape, strides, padding[, ...])</code>	Max pool.

Average Pool

`haiku.avg_pool(value, window_shape, strides, padding, channel_axis=-1)`

Average pool.

Parameters

- **value** (*jax.Array*) – Value to pool.
- **window_shape** (*Union[int, Sequence[int]]*) – Shape of the pooling window, same rank as `value`.
- **strides** (*Union[int, Sequence[int]]*) – Strides of the pooling window, same rank as `value`.
- **padding** (*str*) – Padding algorithm. Either `VALID` or `SAME`.
- **channel_axis** (*Optional[int]*) – Axis of the spatial channels for which pooling is skipped.

Return type `jax.Array`

Returns Pooled result. Same rank as `value`.

Raises **ValueError** – If the padding is not valid.

class haiku.**AvgPool**(*window_shape, strides, padding, channel_axis=- 1, name=None*)

Average pool.

Equivalent to partial application of `avg_pool()`.

__init__(*window_shape, strides, padding, channel_axis=- 1, name=None*)

Average pool.

Parameters

- **window_shape** (*Union[int, Sequence[int]]*) – Shape of the pooling window, same rank as value.
- **strides** (*Union[int, Sequence[int]]*) – Strides of the pooling window, same rank as value.
- **padding** (*str*) – Padding algorithm. Either VALID or SAME.
- **channel_axis** (*Optional[int]*) – Axis of the spatial channels for which pooling is skipped.
- **name** (*Optional[str]*) – String name for the module.

__call__(*value*)

Call self as a function.

Return type jax.Array

Max Pool

haiku.**max_pool**(*value, window_shape, strides, padding, channel_axis=- 1*)

Max pool.

Parameters

- **value** (*jax.Array*) – Value to pool.
- **window_shape** (*Union[int, Sequence[int]]*) – Shape of the pooling window, same rank as value.
- **strides** (*Union[int, Sequence[int]]*) – Strides of the pooling window, same rank as value.
- **padding** (*str*) – Padding algorithm. Either VALID or SAME.
- **channel_axis** (*Optional[int]*) – Axis of the spatial channels for which pooling is skipped.

Return type jax.Array

Returns Pooled result. Same rank as value.

class haiku.**MaxPool**(*window_shape, strides, padding, channel_axis=- 1, name=None*)

Max pool.

Equivalent to partial application of `max_pool()`.

__init__(*window_shape, strides, padding, channel_axis=- 1, name=None*)

Max pool.

Parameters

- **window_shape** (*Union[int, Sequence[int]]*) – Shape of the pooling window, same rank as value.
- **strides** (*Union[int, Sequence[int]]*) – Strides of the pooling window, same rank as value.
- **padding** (*str*) – Padding algorithm. Either VALID or SAME.
- **channel_axis** (*Optional[int]*) – Axis of the spatial channels for which pooling is skipped.
- **name** (*Optional[str]*) – String name for the module.

`__call__`(*value*)

Call self as a function.

Return type `jax.Array`

Dropout

`dropout`(*rng, rate, x[, broadcast_dims]*)

Randomly drop units in the input at a given rate.

dropout

`haiku.dropout`(*rng, rate, x, broadcast_dims=()*)

Randomly drop units in the input at a given rate.

See: <http://www.cs.toronto.edu/~hinton/absps/dropout.pdf>

Parameters

- **rng** (*PRNGKey*) – A JAX random key.
- **rate** (*float*) – Probability that each element of `x` is discarded. Must be a scalar in the range `[0, 1)`.
- **x** (*jax.Array*) – The value to be dropped out.
- **broadcast_dims** (*Sequence[int]*) – specifies dimensions that will share the same dropout mask.

Return type `jax.Array`

Returns `x`, but dropped out and scaled by $1 / (1 - \text{rate})$.

Note: This involves generating `x.size` pseudo-random samples from $U([0, 1))$ computed with the full precision required to compare them with `rate`. When `rate` is a Python float, this is typically 32 bits, which is often more than what applications require. A work-around is to pass `rate` with a lower precision, e.g. using `np.float16(rate)`.

Combinator

<code>Sequential(layers[, name])</code>	Sequentially calls the given list of layers.
---	--

Sequential

class `haiku.Sequential(layers, name=None)`

Sequentially calls the given list of layers.

Note that `Sequential` is limited in the range of possible architectures it can handle. This is a deliberate design decision; `Sequential` is only meant to be used for the simple case of fusing together modules/ops where the input of a particular module/op is the output of the previous one.

Another restriction is that it is not possible to have extra arguments in the `__call__()` method that are passed to the constituents of the module - for example, if there is a `BatchNorm` module in `Sequential` and the user wishes to switch the `is_training` flag. If this is the desired use case, the recommended solution is to subclass `Module` and implement `__call__`:

```
>>> class CustomModule(hk.Module):
...     def __call__(self, x, is_training):
...         x = hk.Conv2D(32, 4, 2)(x)
...         x = hk.BatchNorm(True, True, 0.9)(x, is_training)
...         x = jax.nn.relu(x)
...         return x
```

`__init__(layers, name=None)`

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters `name` (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If `name` is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

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

Calls all layers sequentially.

Convolutional

<code>ConvND(num_spatial_dims, output_channels, ...)</code>	General N-dimensional convolutional.
<code>Conv1D(output_channels, kernel_shape[, ...])</code>	One dimensional convolution.
<code>Conv2D(output_channels, kernel_shape[, ...])</code>	Two dimensional convolution.
<code>Conv3D(output_channels, kernel_shape[, ...])</code>	Three dimensional convolution.
<code>ConvNDTranspose(num_spatial_dims, ..., ...)</code>	General n-dimensional transposed convolution (aka.
<code>Conv1DTranspose(output_channels, kernel_shape)</code>	One dimensional transposed convolution (aka.
<code>Conv2DTranspose(output_channels, kernel_shape)</code>	Two dimensional transposed convolution (aka.
<code>Conv3DTranspose(output_channels, kernel_shape)</code>	Three dimensional transposed convolution (aka.
<code>DepthwiseConv1D(channel_multiplier, kernel_shape)</code>	One dimensional convolution.
<code>DepthwiseConv2D(channel_multiplier, kernel_shape)</code>	Two dimensional convolution.
<code>DepthwiseConv3D(channel_multiplier, kernel_shape)</code>	Three dimensional convolution.
<code>get_channel_index(data_format)</code>	Returns the channel index when given a valid data format.

ConvND

```
class haiku.ConvND(num_spatial_dims, output_channels, kernel_shape, stride=1, rate=1, padding='SAME',
                  with_bias=True, w_init=None, b_init=None, data_format='channels_last', mask=None,
                  feature_group_count=1, name=None)
```

General N-dimensional convolutional.

```
__init__(num_spatial_dims, output_channels, kernel_shape, stride=1, rate=1, padding='SAME',
         with_bias=True, w_init=None, b_init=None, data_format='channels_last', mask=None,
         feature_group_count=1, name=None)
```

Initializes the module.

Parameters

- **num_spatial_dims** (*int*) – The number of spatial dimensions of the input.
- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length `num_spatial_dims`.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length `num_spatial_dims`. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length `num_spatial_dims`. 1 corresponds to standard ND convolution, `rate > 1` corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]], hk.pad.PadFn, Sequence[hk.pad.PadFn]]*) – Optional padding algorithm. Either `VALID` or `SAME` or a sequence of `n` (`low`, `high`) integer pairs that give the padding to apply before and after each spatial dimension. or a callable or sequence of callables of size `num_spatial_dims`. Any callables must take a single integer argument equal to the effective kernel size and return a sequence of two integers representing the padding before and after. See `haiku.pad.*` for more details and example functions. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, `true`.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default, `channels_last`. See `get_channel_index()`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **feature_group_count** (*int*) – Optional number of groups in group convolution. Default value of 1 corresponds to normal dense convolution. If a higher value is used, convolutions are applied separately to that many groups, then stacked together. This reduces the number of parameters and possibly the compute for a given `output_channels`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **name** (*Optional[str]*) – The name of the module.

`__call__(inputs, *, precision=None)`

Connects ConvND layer.

Parameters

- **inputs** (*jax.Array*) – An array of shape `[spatial_dims, C]` and rank- $N+1$ if unbatched, or an array of shape `[N, spatial_dims, C]` and rank- $N+2$ if batched.
- **precision** (*Optional[lax.Precision]*) – Optional `jax.lax.Precision` to pass to `jax.lax.conv_general_dilated()`.

Return type `jax.Array`

Returns

An array of shape `[spatial_dims, output_channels]` and rank- $N+1$ if unbatched, or an array of shape `[N, spatial_dims, output_channels]` and rank- $N+2$ if batched.

Conv1D

```
class haiku.Conv1D(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
                  w_init=None, b_init=None, data_format='NWC', mask=None, feature_group_count=1,
                  name=None)
```

One dimensional convolution.

```
__init__(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
         w_init=None, b_init=None, data_format='NWC', mask=None, feature_group_count=1,
         name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 1.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 1. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 1. 1 corresponds to standard ND convolution, $rate > 1$ corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]], hk.pad.PadFn, Sequence[hk.pad.PadFn]]*) – Optional padding algorithm. Either `VALID` or `SAME` or a callable or sequence of callables of length 1. Any callables must take a single integer argument equal to the effective kernel size and return a list of two integers representing the padding before and after. See `haiku.pad.*` for more details and example functions. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either `NWC` or `NCW`. By default, `NWC`.

- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **feature_group_count** (*int*) – Optional number of groups in group convolution. Default value of 1 corresponds to normal dense convolution. If a higher value is used, convolutions are applied separately to that many groups, then stacked together. This reduces the number of parameters and possibly the compute for a given `output_channels`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **name** (*Optional[str]*) – The name of the module.

Conv2D

```
class haiku.Conv2D(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
                  w_init=None, b_init=None, data_format='NHWC', mask=None, feature_group_count=1,
                  name=None)
```

Two dimensional convolution.

```
__init__(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
         w_init=None, b_init=None, data_format='NHWC', mask=None, feature_group_count=1,
         name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 2.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 2. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 2. 1 corresponds to standard ND convolution, `rate > 1` corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]], hk.pad.PadFn, Sequence[hk.pad.PadFn]]*) – Optional padding algorithm. Either `VALID` or `SAME` or a callable or sequence of callables of length 2. Any callables must take a single integer argument equal to the effective kernel size and return a list of two integers representing the padding before and after. See `haiku.pad.*` for more details and example functions. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, `true`.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either `NHWC` or `NCHW`. By default, `NHWC`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **feature_group_count** (*int*) – Optional number of groups in group convolution. Default value of 1 corresponds to normal dense convolution. If a higher value is used, convolutions are applied separately to that many groups, then stacked together. This reduces the number

of parameters and possibly the compute for a given `output_channels`s. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.

- **name** (*Optional[str]*) – The name of the module.

Conv3D

```
class haiku.Conv3D(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
                  w_init=None, b_init=None, data_format='NDHWC', mask=None, feature_group_count=1,
                  name=None)
```

Three dimensional convolution.

```
__init__(output_channels, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
          w_init=None, b_init=None, data_format='NDHWC', mask=None, feature_group_count=1,
          name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 3.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 3. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 3. 1 corresponds to standard ND convolution, *rate > 1* corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]], hk.pad.PadFn, Sequence[hk.pad.PadFn]]*) – Optional padding algorithm. Either `VALID` or `SAME` or a callable or sequence of callables of length 3. Any callables must take a single integer argument equal to the effective kernel size and return a list of two integers representing the padding before and after. See `haiku.pad.*` for more details and example functions. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either `NDHWC` or `NCDHW`. By default, `NDHWC`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **feature_group_count** (*int*) – Optional number of groups in group convolution. Default value of 1 corresponds to normal dense convolution. If a higher value is used, convolutions are applied separately to that many groups, then stacked together. This reduces the number of parameters and possibly the compute for a given `output_channels`s. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **name** (*Optional[str]*) – The name of the module.

ConvNDTranspose

```
class haiku.ConvNDTranspose(num_spatial_dims, output_channels, kernel_shape, stride=1,
                           output_shape=None, padding='SAME', with_bias=True, w_init=None,
                           b_init=None, data_format='channels_last', mask=None, name=None)
```

General n-dimensional transposed convolution (aka. deconvolution).

```
__init__(num_spatial_dims, output_channels, kernel_shape, stride=1, output_shape=None,
         padding='SAME', with_bias=True, w_init=None, b_init=None, data_format='channels_last',
         mask=None, name=None)
```

Initializes the module.

Parameters

- **num_spatial_dims** (*int*) – The number of spatial dimensions of the input.
- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length `num_spatial_dims`.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length `num_spatial_dims`. Defaults to 1.
- **output_shape** (*Optional[Union[int, Sequence[int]]]*) – Output shape of the spatial dimensions of a transpose convolution. Can be either an integer or an iterable of integers. If a *None* value is given, a default shape is automatically calculated.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either “VALID” or “SAME”. Defaults to “SAME”. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default, `channels_last`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **name** (*Optional[str]*) – The name of the module.

```
__call__(inputs, *, precision=None)
```

Computes the transposed convolution of the input.

Parameters

- **inputs** (*jax.Array*) – An array of shape `[spatial_dims, C]` and rank-`N+1` if unbatched, or an array of shape `[N, spatial_dims, C]` and rank-`N+2` if batched.
- **precision** (*Optional[lax.Precision]*) – Optional `jax.lax.Precision` to pass to `jax.lax.conv_transpose()`.

Return type `jax.Array`

Returns

An array of shape `[spatial_dims, output_channels]` and rank- $N+1$ if unbatched, or an array of shape `[N, spatial_dims, output_channels]` and rank- $N+2$ if batched.

Conv1DTranspose

```
class haiku.Conv1DTranspose(output_channels, kernel_shape, stride=1, output_shape=None,
                           padding='SAME', with_bias=True, w_init=None, b_init=None,
                           data_format='NWC', mask=None, name=None)
```

One dimensional transposed convolution (aka. deconvolution).

```
__init__(output_channels, kernel_shape, stride=1, output_shape=None, padding='SAME', with_bias=True,
         w_init=None, b_init=None, data_format='NWC', mask=None, name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 1.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 1. Defaults to 1.
- **output_shape** (*Optional[Union[int, Sequence[int]]]*) – Output shape of the spatial dimensions of a transpose convolution. Can be either an integer or an iterable of integers. If a *None* value is given, a default shape is automatically calculated.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either `VALID` or `SAME`. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either `NWC` or `NCW`. By default, `NWC`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **name** (*Optional[str]*) – The name of the module.

Conv2DTranspose

```
class haiku.Conv2DTranspose(output_channels, kernel_shape, stride=1, output_shape=None,
                           padding='SAME', with_bias=True, w_init=None, b_init=None,
                           data_format='NHWC', mask=None, name=None)
```

Two dimensional transposed convolution (aka. deconvolution).

```
__init__(output_channels, kernel_shape, stride=1, output_shape=None, padding='SAME', with_bias=True,
         w_init=None, b_init=None, data_format='NHWC', mask=None, name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 2.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 2. Defaults to 1.
- **output_shape** (*Optional[Union[int, Sequence[int]]]*) – Output shape of the spatial dimensions of a transpose convolution. Can be either an integer or an iterable of integers. If a *None* value is given, a default shape is automatically calculated.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either `VALID` or `SAME`. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either `NHWC` or `NCHW`. By default, `NHWC`.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **name** (*Optional[str]*) – The name of the module.

Conv3DTranspose

```
class haiku.Conv3DTranspose(output_channels, kernel_shape, stride=1, output_shape=None,
                           padding='SAME', with_bias=True, w_init=None, b_init=None,
                           data_format='NDHWC', mask=None, name=None)
```

Three dimensional transposed convolution (aka. deconvolution).

```
__init__(output_channels, kernel_shape, stride=1, output_shape=None, padding='SAME', with_bias=True,
          w_init=None, b_init=None, data_format='NDHWC', mask=None, name=None)
```

Initializes the module.

Parameters

- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 3.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 3. Defaults to 1.
- **output_shape** (*Optional[Union[int, Sequence[int]]]*) – Output shape of the spatial dimensions of a transpose convolution. Can be either an integer or an iterable of integers. If a *None* value is given, a default shape is automatically calculated.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either `VALID` or `SAME`. Defaults to `SAME`. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.

- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Either NDHWC or NCDHW. By default, NDHWC.
- **mask** (*Optional[jax.Array]*) – Optional mask of the weights.
- **name** (*Optional[str]*) – The name of the module.

DepthwiseConv1D

```
class haiku.DepthwiseConv1D(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME',
                             with_bias=True, w_init=None, b_init=None, data_format='NWC', name=None)
```

One dimensional convolution.

```
__init__(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
          w_init=None, b_init=None, data_format='NWC', name=None)
```

Construct a 1D Depthwise Convolution.

Parameters

- **channel_multiplier** (*int*) – Multiplicity of output channels. To keep the number of output channels the same as the number of input channels, set 1.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 1.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 1. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 1. 1 corresponds to standard ND convolution, $rate > 1$ corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either VALID, SAME or a sequence of before, after pairs. Defaults to SAME. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Can be either channels_first, channels_last, N...C or NC.... By default, channels_last. See [get_channel_index\(\)](#).
- **name** (*Optional[str]*) – The name of the module.

DepthwiseConv2D

```
class haiku.DepthwiseConv2D(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME',
                             with_bias=True, w_init=None, b_init=None, data_format='NHWC',
                             name=None)
```

Two dimensional convolution.

```
__init__(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
          w_init=None, b_init=None, data_format='NHWC', name=None)
```

Construct a 2D Depthwise Convolution.

Parameters

- **channel_multiplier** (*int*) – Multiplicity of output channels. To keep the number of output channels the same as the number of input channels, set 1.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 2.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 2. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 1. 1 corresponds to standard ND convolution, `rate > 1` corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either VALID, SAME or a sequence of before, after pairs. Defaults to SAME. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Can be either channels_first, channels_last, N...C or NC... By default, channels_last. See [get_channel_index\(\)](#).
- **name** (*Optional[str]*) – The name of the module.

DepthwiseConv3D

```
class haiku.DepthwiseConv3D(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME',
                             with_bias=True, w_init=None, b_init=None, data_format='NDHWC',
                             name=None)
```

Three dimensional convolution.

```
__init__(channel_multiplier, kernel_shape, stride=1, rate=1, padding='SAME', with_bias=True,
          w_init=None, b_init=None, data_format='NDHWC', name=None)
```

Construct a 3D Depthwise Convolution.

Parameters

- **channel_multiplier** (*int*) – Multiplicity of output channels. To keep the number of output channels the same as the number of input channels, set 1.

- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length 3.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length 3. Defaults to 1.
- **rate** (*Union[int, Sequence[int]]*) – Optional kernel dilation rate. Either an integer or a sequence of length 1. 1 corresponds to standard ND convolution, `rate > 1` corresponds to dilated convolution. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either VALID, SAME or a sequence of before, after pairs. Defaults to SAME. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.
- **data_format** (*str*) – The data format of the input. Can be either channels_first, channels_last, N...C or NC.... By default, channels_last. See [get_channel_index\(\)](#).
- **name** (*Optional[str]*) – The name of the module.

SeparableDepthwiseConv2D

```
class haiku.SeparableDepthwiseConv2D(channel_multiplier, kernel_shape, stride=1, padding='SAME',
                                     with_bias=True, w_init=None, b_init=None, data_format='NHWC',
                                     name=None)
```

Separable 2-D Depthwise Convolution Module.

```
__init__(channel_multiplier, kernel_shape, stride=1, padding='SAME', with_bias=True, w_init=None,
         b_init=None, data_format='NHWC', name=None)
```

Construct a Separable 2D Depthwise Convolution module.

Parameters

- **channel_multiplier** (*int*) – Multiplicity of output channels. To keep the number of output channels the same as the number of input channels, set 1.
- **kernel_shape** (*Union[int, Sequence[int]]*) – The shape of the kernel. Either an integer or a sequence of length `num_spatial_dims`.
- **stride** (*Union[int, Sequence[int]]*) – Optional stride for the kernel. Either an integer or a sequence of length `num_spatial_dims`. Defaults to 1.
- **padding** (*Union[str, Sequence[tuple[int, int]]]*) – Optional padding algorithm. Either VALID, SAME or a sequence of before, after pairs. Defaults to SAME. See: https://www.tensorflow.org/xla/operation_semantics#conv_convolution.
- **with_bias** (*bool*) – Whether to add a bias. By default, true.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Optional weight initialization. By default, truncated normal.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Optional bias initialization. By default, zeros.

- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default, `channels_last`.
- **name** (*Optional[str]*) – The name of the module.

`__call__` (*inputs*)

Call self as a function.

Return type `jax.Array`

get_channel_index

`haiku.get_channel_index` (*data_format*)

Returns the channel index when given a valid data format.

```
>>> hk.get_channel_index('channels_last')
-1
>>> hk.get_channel_index('channels_first')
1
>>> hk.get_channel_index('N...C')
-1
>>> hk.get_channel_index('NCHW')
1
```

Parameters `data_format` (*str*) – String, the data format to get the channel index from. Valid data formats are spatial (e.g. ```NCHW```), sequential (e.g. `BTHWD`), `channels_first` and `channels_last`).

Return type `int`

Returns The channel index as an `int`, either 1 or -1.

Raises `ValueError` – If the data format is unrecognised.

Normalization

<code>BatchNorm</code> (<i>create_scale, create_offset, ...</i>)	Normalizes inputs to maintain a mean of ~0 and stddev of ~1.
<code>GroupNorm</code> (<i>groups[, axis, create_scale, ...]</i>)	Group normalization module.
<code>InstanceNorm</code> (<i>create_scale, create_offset[, ...]</i>)	Normalizes inputs along the spatial dimensions.
<code>LayerNorm</code> (<i>axis, create_scale, create_offset</i>)	LayerNorm module.
<code>RMSNorm</code> (<i>axis[, eps, scale_init, name, ...]</i>)	RMSNorm module.
<code>SpectralNorm</code> (<i>[eps, n_steps, name]</i>)	Normalizes an input by its first singular value.
<code>ExponentialMovingAverage</code> (<i>decay[, ...]</i>)	Maintains an exponential moving average.
<code>SNParamsTree</code> (<i>[eps, n_steps, ignore_regex, name]</i>)	Applies Spectral Normalization to all parameters in a tree.
<code>EMAPParamsTree</code> (<i>decay[, zero_debias, ...]</i>)	Maintains an exponential moving average for all parameters in a tree.

BatchNorm

```
class haiku.BatchNorm(create_scale, create_offset, decay_rate, eps=1e-05, scale_init=None, offset_init=None,
                       axis=None, cross_replica_axis=None, cross_replica_axis_index_groups=None,
                       data_format='channels_last', name=None)
```

Normalizes inputs to maintain a mean of ~ 0 and stddev of ~ 1 .

See: <https://arxiv.org/abs/1502.03167>.

There are many different variations for how users want to manage scale and offset if they require them at all. These are:

- No scale/offset in which case `create_*` should be set to `False` and `scale/offset` aren't passed when the module is called.
- Trainable scale/offset in which case `create_*` should be set to `True` and again `scale/offset` aren't passed when the module is called. In this case this module creates and owns the `scale/offset` variables.
- Externally generated `scale/offset`, such as for conditional normalization, in which case `create_*` should be set to `False` and then the values fed in at call time.

NOTE: `jax.vmap(hk.transform(BatchNorm))` will update summary statistics and normalize values on a per-batch basis; we currently do *not* support normalizing across a batch axis introduced by `vmap`.

```
__init__(create_scale, create_offset, decay_rate, eps=1e-05, scale_init=None, offset_init=None, axis=None,
          cross_replica_axis=None, cross_replica_axis_index_groups=None, data_format='channels_last',
          name=None)
```

Constructs a BatchNorm module.

Parameters

- **create_scale** (*bool*) – Whether to include a trainable scaling factor.
- **create_offset** (*bool*) – Whether to include a trainable offset.
- **decay_rate** (*float*) – Decay rate for EMA.
- **eps** (*float*) – Small epsilon to avoid division by zero variance. Defaults $1e-5$, as in the paper and Sonnet.
- **scale_init** (*Optional* [*hk.initializers.Initializer*]) – Optional initializer for gain (aka scale). Can only be set if `create_scale=True`. By default, `1`.
- **offset_init** (*Optional* [*hk.initializers.Initializer*]) – Optional initializer for bias (aka offset). Can only be set if `create_offset=True`. By default, `0`.
- **axis** (*Optional* [*Sequence* [*int*]]) – Which axes to reduce over. The default (`None`) signifies that all but the channel axis should be normalized. Otherwise this is a list of axis indices which will have normalization statistics calculated.
- **cross_replica_axis** (*Optional* [*Union* [*str*, *Sequence* [*str*]]) – If not `None`, it should be a string (or sequence of strings) representing the axis name(s) over which this module is being run within a `jax` map (e.g. `jax.pmap` or `jax.vmap`). Supplying this argument means that batch statistics are calculated across all replicas on the named axes.
- **cross_replica_axis_index_groups** (*Optional* [*Sequence* [*Sequence* [*int*]]) – Specifies how devices are grouped. Valid only within `jax.pmap` collectives.
- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default it is `channels_last`. See `get_channel_index()`.

- **name** (*Optional[str]*) – The module name.

`__call__(inputs, is_training, test_local_stats=False, scale=None, offset=None)`

Computes the normalized version of the input.

Parameters

- **inputs** (*jax.Array*) – An array, where the data format is [..., C].
- **is_training** (*bool*) – Whether this is during training.
- **test_local_stats** (*bool*) – Whether local stats are used when `is_training=False`.
- **scale** (*Optional[jax.Array]*) – An array up to n-D. The shape of this tensor must be broadcastable to the shape of `inputs`. This is the scale applied to the normalized inputs. This cannot be passed in if the module was constructed with `create_scale=True`.
- **offset** (*Optional[jax.Array]*) – An array up to n-D. The shape of this tensor must be broadcastable to the shape of `inputs`. This is the offset applied to the normalized inputs. This cannot be passed in if the module was constructed with `create_offset=True`.

Return type `jax.Array`

Returns The array, normalized across all but the last dimension.

GroupNorm

```
class haiku.GroupNorm(groups, axis=slice(1, None, None), create_scale=True, create_offset=True, eps=1e-05,
                      scale_init=None, offset_init=None, data_format='channels_last', name=None)
```

Group normalization module.

This applies group normalization to the `x`. This involves splitting the channels into groups before calculating the mean and variance. The default behaviour is to compute the mean and variance over the spatial dimensions and the grouped channels. The mean and variance will never be computed over the created groups axis.

It transforms the input `x` into:

$$\text{outputs} = \text{scale} \frac{x - \mu}{\sigma + \epsilon} + \text{offset}$$

Where μ and σ are respectively the mean and standard deviation of `x`.

There are many different variations for how users want to manage scale and offset if they require them at all. These are:

- No scale/offset in which case `create_*` should be set to `False` and `scale/offset` aren't passed when the module is called.
- Trainable scale/offset in which case `create_*` should be set to `True` and again `scale/offset` aren't passed when the module is called. In this case this module creates and owns the `scale/offset` parameters.
- Externally generated scale/offset, such as for conditional normalization, in which case `create_*` should be set to `False` and then the values fed in at call time.

```
__init__(groups, axis=slice(1, None, None), create_scale=True, create_offset=True, eps=1e-05,
         scale_init=None, offset_init=None, data_format='channels_last', name=None)
```

Constructs a `GroupNorm` module.

Parameters

- **groups** (*int*) – number of groups to divide the channels by. The number of channels must be divisible by this.

- **axis** (*Union[int, slice, Sequence[int]]*) – int, slice or sequence of ints representing the axes which should be normalized across. By default this is all but the first dimension. For time series data use *slice(2, None)* to average over the none Batch and Time data.
- **create_scale** (*bool*) – whether to create a trainable scale per channel applied after the normalization.
- **create_offset** (*bool*) – whether to create a trainable offset per channel applied after normalization and scaling.
- **eps** (*float*) – Small epsilon to add to the variance to avoid division by zero. Defaults to 1e-5.
- **scale_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for the scale parameter. Can only be set if `create_scale=True`. By default scale is initialized to 1.
- **offset_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for the offset parameter. Can only be set if `create_offset=True`. By default offset is initialized to 0.
- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default it is `channels_last`. See [get_channel_index\(\)](#).
- **name** (*Optional[str]*) – Name of the module.

`__call__(x, scale=None, offset=None)`

Returns normalized inputs.

Parameters

- **x** (*jax.Array*) – An n-D tensor of the `data_format` specified in the constructor on which the transformation is performed.
- **scale** (*Optional[jax.Array]*) – A tensor up to n-D. The shape of this tensor must be broadcastable to the shape of `x`. This is the scale applied to the normalized `x`. This cannot be passed in if the module was constructed with `create_scale=True`.
- **offset** (*Optional[jax.Array]*) – A tensor up to n-D. The shape of this tensor must be broadcastable to the shape of `x`. This is the offset applied to the normalized `x`. This cannot be passed in if the module was constructed with `create_offset=True`.

Return type `jax.Array`

Returns An n-d tensor of the same shape as `x` that has been normalized.

InstanceNorm

```
class haiku.InstanceNorm(create_scale, create_offset, eps=1e-05, scale_init=None, offset_init=None,
                        data_format='channels_last', name=None)
```

Normalizes inputs along the spatial dimensions.

See [LayerNorm](#) for more details.

```
__init__(create_scale, create_offset, eps=1e-05, scale_init=None, offset_init=None,
         data_format='channels_last', name=None)
```

Constructs an [InstanceNorm](#) module.

This method creates a module which normalizes over the spatial dimensions.

Parameters

- **create_scale** (*bool*) – *bool* representing whether to create a trainable scale per channel applied after the normalization.
- **create_offset** (*bool*) – *bool* representing whether to create a trainable offset per channel applied after normalization and scaling.
- **eps** (*float*) – Small epsilon to avoid division by zero variance. Defaults to $1e-5$.
- **scale_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for the scale variable. Can only be set if `create_scale=True`. By default scale is initialized to 1.
- **offset_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for the offset variable. Can only be set if `create_offset=True`. By default offset is initialized to 0.
- **data_format** (*str*) – The data format of the input. Can be either `channels_first`, `channels_last`, `N...C` or `NC...`. By default it is `channels_last`. See `get_channel_index()`.
- **name** (*Optional[str]*) – Name of the module.

LayerNorm

```
class haiku.LayerNorm(axis, create_scale, create_offset, eps=1e-05, scale_init=None, offset_init=None,
                      use_fast_variance=False, name=None, *, param_axis=None)
```

LayerNorm module.

See: <https://arxiv.org/abs/1607.06450>.

Example usage:

```
>>> ln = hk.LayerNorm(axis=-1, param_axis=-1,
...                   create_scale=True, create_offset=True)
>>> x = ln(jnp.ones([8, 224, 224, 3]))
```

```
__init__(axis, create_scale, create_offset, eps=1e-05, scale_init=None, offset_init=None,
         use_fast_variance=False, name=None, *, param_axis=None)
```

Constructs a LayerNorm module.

Parameters

- **axis** (*AxisOrAxes*) – Integer, list of integers, or slice indicating which axes to normalize over. Note that the shape of the scale/offset parameters are controlled by the `param_axis` argument.
- **create_scale** (*bool*) – *Bool*, defines whether to create a trainable scale per channel applied after the normalization.
- **create_offset** (*bool*) – *Bool*, defines whether to create a trainable offset per channel applied after normalization and scaling.
- **eps** (*float*) – Small epsilon to avoid division by zero variance. Defaults $1e-5$, as in the paper and Sonnet.
- **scale_init** (*Optional[hk.initializers.Initializer]*) – Optional initializer for gain (aka scale). By default, one.

- **offset_init** (*Optional*[*hk.initializers.Initializer*]) – Optional initializer for bias (aka offset). By default, zero.
- **use_fast_variance** (*bool*) – If true, use a faster but less numerically stable formulation for computing variance.
- **name** (*Optional*[*str*]) – The module name.
- **param_axis** (*Optional*[*AxisOrAxes*]) – Axis used to determine the parameter shape of the learnable scale/offset. Sonnet sets this to the channel/feature axis (e.g. to -1 for NHWC). Other libraries set this to the same as the reduction axis (e.g. `axis=param_axis`).

`__call__` (*inputs*, *scale=None*, *offset=None*)

Connects the layer norm.

Parameters

- **inputs** (*jax.Array*) – An array, where the data format is [N, ..., C].
- **scale** (*Optional*[*jax.Array*]) – An array up to n-D. The shape of this tensor must be broadcastable to the shape of `inputs`. This is the scale applied to the normalized inputs. This cannot be passed in if the module was constructed with `create_scale=True`.
- **offset** (*Optional*[*jax.Array*]) – An array up to n-D. The shape of this tensor must be broadcastable to the shape of `inputs`. This is the offset applied to the normalized inputs. This cannot be passed in if the module was constructed with `create_offset=True`.

Return type *jax.Array*

Returns The array, normalized.

RMSNorm

`class haiku.RMSNorm` (*axis*, *eps=1e-05*, *scale_init=None*, *name=None*, *create_scale=True*, *, *param_axis=None*)

RMSNorm module.

RMSNorm provides an alternative that can be both faster and more stable than LayerNorm. The inputs are normalized by the root-mean-squared (RMS) and scaled by a learned parameter, but they are not recentered around their mean.

See <https://arxiv.org/pdf/1910.07467.pdf>

`__init__` (*axis*, *eps=1e-05*, *scale_init=None*, *name=None*, *create_scale=True*, *, *param_axis=None*)

Constructs a RMSNorm module.

Parameters

- **axis** (*AxisOrAxes*) – Integer, list of integers, or slice indicating which axes to normalize over.
- **eps** (*float*) – Small epsilon to avoid division by zero variance. Defaults to 1e-5.
- **scale_init** (*Optional*[*hk.initializers.Initializer*]) – Optional initializer for gain (aka scale). By default, one.
- **name** (*Optional*[*str*]) – The module name.
- **create_scale** (*bool*) – Bool, defines whether to create a trainable scale per channel applied after the normalization.

- **param_axis** (*Optional*[*AxisOrAxes*]) – Axis used to determine the parameter shape of the learnable scale/offset. Sonnet sets this to the channel/feature axis (e.g. to -1 for NHWC). Other libraries set this to the same as the reduction axis (e.g. `axis=param_axis`). *None* defaults to (-1).

`__call__(inputs)`

Connects the layer norm.

Parameters `inputs` (*jax.Array*) – An array, where the data format is [N, . . . , C].

Returns The normalized array, of the same shape as the inputs.

SpectralNorm

`class haiku.SpectralNorm(eps=0.0001, n_steps=1, name=None)`

Normalizes an input by its first singular value.

This module uses power iteration to calculate this value based on the input and an internal hidden state.

`__init__(eps=0.0001, n_steps=1, name=None)`

Initializes an SpectralNorm module.

Parameters

- **eps** (*float*) – The constant used for numerical stability.
- **n_steps** (*int*) – How many steps of power iteration to perform to approximate the singular value of the input.
- **name** (*Optional*[*str*]) – The name of the module.

`__call__(value, update_stats=True, error_on_non_matrix=False)`

Performs Spectral Normalization and returns the new value.

Parameters

- **value** – The array-like object for which you would like to perform an spectral normalization on.
- **update_stats** (*bool*) – A boolean defaulting to True. Regardless of this arg, this function will return the normalized input. When `update_stats` is True, the internal state of this object will also be updated to reflect the input value. When `update_stats` is False the internal stats will remain unchanged.
- **error_on_non_matrix** (*bool*) – Spectral normalization is only defined on matrices. By default, this module will return scalars unchanged and flatten higher-order tensors in their leading dimensions. Setting this flag to True will instead throw errors in those cases.

Return type `jax.Array`

Returns The input value normalized by it's first singular value.

Raises **ValueError** – If `error_on_non_matrix` is True and `value` has `ndims > 2`.

ExponentialMovingAverage

class haiku.ExponentialMovingAverage(*decay*, *zero_debias=True*, *warmup_length=0*, *name=None*)

Maintains an exponential moving average.

This uses the Adam debiasing procedure. See <https://arxiv.org/pdf/1412.6980.pdf> for details.

__init__(*decay*, *zero_debias=True*, *warmup_length=0*, *name=None*)

Initializes an ExponentialMovingAverage module.

Parameters

- **decay** – The chosen decay. Must in $[0, 1)$. Values close to 1 result in slow decay; values close to 0 result in fast decay.
- **zero_debias** (*bool*) – Whether to run with zero-debiasing.
- **warmup_length** (*int*) – A positive integer, EMA has no effect until the internal counter has reached *warmup_length* at which point the initial value for the decaying average is initialized to the input value after *warmup_length* iterations.
- **name** (*Optional[str]*) – The name of the module.

initialize(*shape*, *dtype=<class 'jax.numpy.float32'>*)

If uninitialized sets the average to zeros of the given shape/dtype.

__call__(*value*, *update_stats=True*)

Updates the EMA and returns the new value.

Parameters

- **value** (*Union[float, jax.Array]*) – The array-like object for which you would like to perform an exponential decay on.
- **update_stats** (*bool*) – A Boolean, whether to update the internal state of this object to reflect the input value. When *update_stats* is False the internal stats will remain unchanged.

Return type *jax.Array*

Returns The exponentially weighted average of the input value.

SNParamsTree

class haiku.SNParamsTree(*eps=0.0001*, *n_steps=1*, *ignore_regex=""*, *name=None*)

Applies Spectral Normalization to all parameters in a tree.

This is isomorphic to EMAParamsTree in *moving_averages.py*.

__init__(*eps=0.0001*, *n_steps=1*, *ignore_regex=""*, *name=None*)

Initializes an SNParamsTree module.

Parameters

- **eps** (*float*) – The constant used for numerical stability.
- **n_steps** (*int*) – How many steps of power iteration to perform to approximate the singular value of the input.
- **ignore_regex** (*str*) – A string. Any parameter in the tree whose name matches this regex will not have spectral normalization applied to it. The empty string means this module applies to all parameters.

- **name** (*Optional[str]*) – The name of the module.

`__call__(tree, update_stats=True)`

Call self as a function.

EMAParamsTree

class haiku.EMAParamsTree(*decay, zero_debias=True, warmup_length=0, ignore_regex="", name=None*)

Maintains an exponential moving average for all parameters in a tree.

While ExponentialMovingAverage is meant to be applied to single parameters within a function, this class is meant to be applied to the entire tree of parameters for a function.

Given a set of parameters for some network:

```
>>> network_fn = lambda x: hk.Linear(10)(x)
>>> x = jnp.ones([1, 1])
>>> params = hk.transform(network_fn).init(jax.random.PRNGKey(428), x)
```

You might use the EMAParamsTree like follows:

```
>>> ema_fn = hk.transform_with_state(lambda x: hk.EMAParamsTree(0.2)(x))
>>> _, ema_state = ema_fn.init(None, params)
>>> ema_params, ema_state = ema_fn.apply(None, ema_state, None, params)
```

Here, we are transforming a Haiku function and constructing its parameters via an `init_fn` as normal, but are creating a second transformed function which expects a tree of parameters as input. This function is then called with the current parameters as input, which then returns an identical tree with every parameter replaced with its exponentially decayed average. This `ema_params` object can then be passed into the `network_fn` as usual, and will cause it to run with EMA weights.

`__init__(decay, zero_debias=True, warmup_length=0, ignore_regex="", name=None)`

Initializes an EMAParamsTree module.

Parameters

- **decay** – The chosen decay. Must in $[0, 1)$. Values close to 1 result in slow decay; values close to 0 result in fast decay.
- **zero_debias** (*bool*) – Whether to run with zero-debiasing.
- **warmup_length** (*int*) – A positive integer, EMA has no effect until the internal counter has reached `warmup_length` at which point the initial value for the decaying average is initialized to the input value after `warmup_length` iterations.
- **ignore_regex** (*str*) – A string. Any parameter in the tree whose name matches this regex will not have any moving average applied to it. The empty string means this module will EMA all parameters.
- **name** (*Optional[str]*) – The name of the module.

`__call__(tree, update_stats=True)`

Call self as a function.

Recurrent

<code>RNNCore([name])</code>	Base class for RNN cores.
<code>dynamic_unroll(core, input_sequence, ..., ...)</code>	Performs a dynamic unroll of an RNN.
<code>static_unroll(core, input_sequence, ..., ...)</code>	Performs a static unroll of an RNN.
<code>expand_apply(f[, axis])</code>	Wraps <code>f</code> to temporarily add a size-1 axis to its inputs.
<code>VanillaRNN(hidden_size[, double_bias, name])</code>	Basic fully-connected RNN core.
<code>LSTM(hidden_size[, name])</code>	Long short-term memory (LSTM) RNN core.
<code>GRU(hidden_size[, w_i_init, w_h_init, ...])</code>	Gated Recurrent Unit.
<code>DeepRNN(layers[, name])</code>	Wraps a sequence of cores and callables as a single core.
<code>deep_rnn_with_skip_connections(layers[, name])</code>	Constructs a <code>DeepRNN</code> with skip connections.
<code>ResetCore(core[, name])</code>	A wrapper for managing state resets during unrolls.
<code>IdentityCore([name])</code>	A recurrent core that forwards the inputs and an empty state.
<code>Conv1DLSTM(input_shape, output_channels, ...)</code>	1-D convolutional LSTM.
<code>Conv2DLSTM(input_shape, output_channels, ...)</code>	2-D convolutional LSTM.
<code>Conv3DLSTM(input_shape, output_channels, ...)</code>	3-D convolutional LSTM.

RNNCore

class `haiku.RNNCore`(*name=None*)

Base class for RNN cores.

This class defines the basic functionality that every core should implement: `initial_state()`, used to construct an example of the core state; and `__call__()` which applies the core parameterized by a previous state to an input.

Cores may be used with `dynamic_unroll()` and `static_unroll()` to iteratively construct an output sequence from the given input sequence.

abstract `__call__`(*inputs, prev_state*)

Run one step of the RNN.

Parameters

- **inputs** – An arbitrarily nested structure.
- **prev_state** – Previous core state.

Return type `tuple[Any, Any]`

Returns A tuple with two elements `output`, `next_state`. `output` is an arbitrarily nested structure. `next_state` is the next core state, this must be the same shape as `prev_state`.

abstract `initial_state`(*batch_size*)

Constructs an initial state for this core.

Parameters `batch_size` (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If `None`, the core may either fail or (experimentally) return an initial state without a batch dimension.

Returns Arbitrarily nested initial state for this core.

dynamic_unroll

`haiku.dynamic_unroll(core, input_sequence, initial_state, time_major=True, reverse=False, return_all_states=False, unroll=1)`

Performs a dynamic unroll of an RNN.

An *unroll* corresponds to calling the core on each element of the input sequence in a loop, carrying the state through:

```
state = initial_state
for t in range(len(input_sequence)):
    outputs, state = core(input_sequence[t], state)
```

A *dynamic* unroll preserves the loop structure when executed inside `jax.jit()`. See `static_unroll()` for an unroll function which replaces a loop with its body repeated multiple times.

Parameters

- **core** – An *RNNCore* to unroll.
- **input_sequence** – An arbitrarily nested structure of tensors of shape `[T, ...]` if `time_major=True`, or `[B, T, ...]` if `time_major=False`, where `T` is the number of time steps.
- **initial_state** – An initial state of the given core.
- **time_major** – If `True`, inputs are expected time-major, otherwise they are expected batch-major.
- **reverse** – If `True`, inputs are scanned in the reversed order. Equivalent to reversing the time dimension in both inputs and outputs. See https://jax.readthedocs.io/en/latest/_autosummary/jax.lax.scan.html for more details.
- **return_all_states** – If `True`, all intermediate states are returned rather than only the last one in time.
- **unroll** – How many scan iterations to unroll within a single iteration of a loop.

Returns

- **output_sequence** - An arbitrarily nested structure of tensors of shape `[T, ...]` if time-major, otherwise `[B, T, ...]`.
- **state_sequence** - If `return_all_states` is `True`, returns the sequence of core states. Otherwise, core state at time step `T`.

Return type A tuple with two elements

static_unroll

`haiku.static_unroll(core, input_sequence, initial_state, time_major=True)`

Performs a static unroll of an RNN.

An *unroll* corresponds to calling the core on each element of the input sequence in a loop, carrying the state through:

```
state = initial_state
for t in range(len(input_sequence)):
    outputs, state = core(input_sequence[t], state)
```

A *static* unroll replaces a loop with its body repeated multiple times when executed inside `jax.jit()`:

```
state = initial_state
outputs0, state = core(input_sequence[0], state)
outputs1, state = core(input_sequence[1], state)
outputs2, state = core(input_sequence[2], state)
...
```

See `dynamic_unroll()` for a loop-preserving unroll function.

Parameters

- **core** – An *RNNCore* to unroll.
- **input_sequence** – An arbitrarily nested structure of tensors of shape `[T, ...]` if `time_major=True`, or `[B, T, ...]` if `time_major=False`, where `T` is the number of time steps.
- **initial_state** – An initial state of the given core.
- **time_major** – If `True`, inputs are expected time-major, otherwise they are expected batch-major.

Returns

- **output_sequence** - An arbitrarily nested structure of tensors of shape `[T, ...]` if time-major, otherwise `[B, T, ...]`.
- **final_state** - Core state at time step `T`.

Return type A tuple with two elements

expand_apply

`haiku.expand_apply(f, axis=0)`

Wraps `f` to temporarily add a size-1 axis to its inputs.

Syntactic sugar for:

```
ins = jax.tree_util.tree_map(lambda t: np.expand_dims(t, axis=axis), ins)
out = f(ins)
out = jax.tree_util.tree_map(lambda t: np.squeeze(t, axis=axis), out)
```

This may be useful for applying a function built for `[Time, Batch, ...]` arrays to a single timestep.

Parameters

- **f** – The callable to be applied to the expanded inputs.
- **axis** – Where to add the extra axis.

Returns `f`, wrapped as described above.

VanillaRNN

class haiku.VanillaRNN(*hidden_size*, *double_bias=True*, *name=None*)

Basic fully-connected RNN core.

Given x_t and the previous hidden state h_{t-1} the core computes

$$h_t = \text{ReLU}(w_i x_t + b_i + w_h h_{t-1} + b_h)$$

The output is equal to the new state, h_t .

__init__(*hidden_size*, *double_bias=True*, *name=None*)

Constructs a vanilla RNN core.

Parameters

- **hidden_size** (*int*) – Hidden layer size.
- **double_bias** (*bool*) – Whether to use a bias in the two linear layers. This changes nothing to the learning performance of the cell. However, doubling will create two sets of bias parameters rather than one.
- **name** (*Optional[str]*) – Name of the module.

__call__(*inputs*, *prev_state*)

Run one step of the RNN.

Parameters

- **inputs** – An arbitrarily nested structure.
- **prev_state** – Previous core state.

Returns A tuple with two elements *output*, *next_state*. *output* is an arbitrarily nested structure. *next_state* is the next core state, this must be the same shape as *prev_state*.

initial_state(*batch_size*)

Constructs an initial state for this core.

Parameters **batch_size** (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If None, the core may either fail or (experimentally) return an initial state without a batch dimension.

Returns Arbitrarily nested initial state for this core.

LSTM

class haiku.LSTM(*hidden_size*, *name=None*)

Long short-term memory (LSTM) RNN core.

The implementation is based on [1]. Given x_t and the previous state (h_{t-1}, c_{t-1}) the core computes

$$\begin{aligned} i_t &= \sigma(W_{ii}x_t + W_{hi}h_{t-1} + b_i) \\ f_t &= \sigma(W_{if}x_t + W_{hf}h_{t-1} + b_f) \\ g_t &= \tanh(W_{ig}x_t + W_{hg}h_{t-1} + b_g) \\ o_t &= \sigma(W_{io}x_t + W_{ho}h_{t-1} + b_o) \\ c_t &= f_t c_{t-1} + i_t g_t \\ h_t &= o_t \tanh(c_t) \end{aligned}$$

where i_t , f_t , o_t are input, forget and output gate activations, and g_t is a vector of cell updates.

The output is equal to the new hidden, h_t .

Notes

Forget gate initialization: Following [2] we add 1.0 to b_f after initialization in order to reduce the scale of forgetting in the beginning of the training.

`__init__(hidden_size, name=None)`

Constructs an LSTM.

Parameters

- **hidden_size** (*int*) – Hidden layer size.
- **name** (*Optional[str]*) – Name of the module.

`__call__(inputs, prev_state)`

Run one step of the RNN.

Parameters

- **inputs** (*jax.Array*) – An arbitrarily nested structure.
- **prev_state** (*LSTMState*) – Previous core state.

Return type *tuple[jax.Array, LSTMState]*

Returns A tuple with two elements `output`, `next_state`. `output` is an arbitrarily nested structure. `next_state` is the next core state, this must be the same shape as `prev_state`.

`initial_state(batch_size)`

Constructs an initial state for this core.

Parameters **batch_size** (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If None, the core may either fail or (experimentally) return an initial state without a batch dimension.

Return type *LSTMState*

Returns Arbitrarily nested initial state for this core.

GRU

`class haiku.GRU(hidden_size, w_i_init=None, w_h_init=None, b_init=None, name=None)`

Gated Recurrent Unit.

The implementation is based on: <https://arxiv.org/pdf/1412.3555v1.pdf> with biases.

Given x_t and the previous state h_{t-1} the core computes

$$\begin{aligned} z_t &= \sigma(W_{iz}x_t + W_{hz}h_{t-1} + b_z) \\ r_t &= \sigma(W_{ir}x_t + W_{hr}h_{t-1} + b_r) \\ a_t &= \tanh(W_{ia}x_t + W_{ha}(r_t \odot h_{t-1}) + b_a) \\ h_t &= (1 - z_t) \odot h_{t-1} + z_t \odot a_t \end{aligned}$$

where z_t and r_t are reset and update gates.

The output is equal to the new hidden state, h_t .

`__init__(hidden_size, w_i_init=None, w_h_init=None, b_init=None, name=None)`

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters *name* (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If *name* is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

`__call__` (*inputs, state*)

Run one step of the RNN.

Parameters

- **inputs** – An arbitrarily nested structure.
- **prev_state** – Previous core state.

Returns A tuple with two elements `output`, `next_state`. `output` is an arbitrarily nested structure. `next_state` is the next core state, this must be the same shape as `prev_state`.

`initial_state` (*batch_size*)

Constructs an initial state for this core.

Parameters *batch_size* (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If None, the core may either fail or (experimentally) return an initial state without a batch dimension.

Returns Arbitrarily nested initial state for this core.

DeepRNN

`class haiku.DeepRNN(layers, name=None)`

Wraps a sequence of cores and callables as a single core.

```
>>> deep_rnn = hk.DeepRNN([
...     hk.LSTM(hidden_size=4),
...     jax.nn.relu,
...     hk.LSTM(hidden_size=2),
... ])
```

The state of a `DeepRNN` is a tuple with one element per `RNNCore`. If no layers are `RNNCores`, the state is an empty tuple.

`__init__` (*layers, name=None*)

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters *name* (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If *name* is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

`haiku.deep_rnn_with_skip_connections(layers, name=None)`

Constructs a `DeepRNN` with skip connections.

Skip connections alter the dependency structure within a `DeepRNN`. Specifically, input to the *i*-th layer (*i* > 0) is given by a concatenation of the core's inputs and the outputs of the (*i*-1)-th layer.

The output of the `DeepRNN` is the concatenation of the outputs of all cores.

```

outputs0, ... = layers[0](inputs, ...)
outputs1, ... = layers[1](tf.concat([inputs, outputs0], axis=-1), ...)
outputs2, ... = layers[2](tf.concat([inputs, outputs1], axis=-1), ...)
...

```

Parameters

- **layers** (*Sequence*[*RNNCore*]) – List of *RNNCores*.
- **name** (*Optional*[*str*]) – Name of the module.

Return type *RNNCore***Returns** A *_DeepRNN* with skip connections.**Raises** **ValueError** – If any of the layers is not an *RNNCore*.**ResetCore****class** `haiku.ResetCore`(*core*, *name=None*)

A wrapper for managing state resets during unrolls.

When unrolling an *RNNCore* on a batch of inputs sequences it may be necessary to reset the core's state at different timesteps for different elements of the batch. The *ResetCore* class enables this by taking a batch of `should_reset` booleans in addition to the batch of inputs, and conditionally resetting the core's state for individual elements of the batch. You may also reset individual entries of the state by passing a `should_reset` nest compatible with the state structure.

__init__(*core*, *name=None*)

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters **name** (*Optional*[*str*]) – An optional string name for the class. Must be a valid Python identifier. If `name` is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

__call__(*inputs*, *state*)

Run one step of the wrapped core, handling state reset.

Parameters

- **inputs** – Tuple with two elements, `inputs`, `should_reset`, where `should_reset` is the signal used to reset the wrapped core's state. `should_reset` can be either tensor or nest. If nest, `should_reset` must match the state structure, and its components' shapes must be prefixes of the corresponding entries tensors' shapes in the state nest. If tensor, supported shapes are all common shape prefixes of the state component tensors, e.g. `[batch_size]`.
- **state** – Previous wrapped core state.

Returns Tuple of the wrapped core's output, `next_state`.**initial_state**(*batch_size*)

Constructs an initial state for this core.

Parameters `batch_size` (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If None, the core may either fail or (experimentally) return an initial state without a batch dimension.

Returns Arbitrarily nested initial state for this core.

IdentityCore

class `haiku.IdentityCore`(*name=None*)

A recurrent core that forwards the inputs and an empty state.

This is commonly used when switching between recurrent and feedforward versions of a model while preserving the same interface.

`__call__`(*inputs, state*)

Run one step of the RNN.

Parameters

- **inputs** – An arbitrarily nested structure.
- **prev_state** – Previous core state.

Returns A tuple with two elements `output`, `next_state`. `output` is an arbitrarily nested structure. `next_state` is the next core state, this must be the same shape as `prev_state`.

`initial_state`(*batch_size*)

Constructs an initial state for this core.

Parameters `batch_size` (*Optional[int]*) – Optional int or an integral scalar tensor representing batch size. If None, the core may either fail or (experimentally) return an initial state without a batch dimension.

Returns Arbitrarily nested initial state for this core.

Conv1DLSTM

class `haiku.Conv1DLSTM`(*input_shape, output_channels, kernel_shape, name=None*)

1-D convolutional LSTM.

The implementation is based on [3]. Given x_t and the previous state (h_{t-1}, c_{t-1}) the core computes

$$\begin{aligned} i_t &= \sigma(W_{ii} * x_t + W_{hi} * h_{t-1} + b_i) \\ f_t &= \sigma(W_{if} * x_t + W_{hf} * h_{t-1} + b_f) \\ g_t &= \tanh(W_{ig} * x_t + W_{hg} * h_{t-1} + b_g) \\ o_t &= \sigma(W_{io} * x_t + W_{ho} * h_{t-1} + b_o) \\ c_t &= f_t c_{t-1} + i_t g_t \\ h_t &= o_t \tanh(c_t) \end{aligned}$$

where $*$ denotes the convolution operator; i_t, f_t, o_t are input, forget and output gate activations, and g_t is a vector of cell updates.

The output is equal to the new hidden state, h_t .

Notes

Forget gate initialization: Following [2] we add 1.0 to b_f after initialization in order to reduce the scale of forgetting in the beginning of the training.

`__init__(input_shape, output_channels, kernel_shape, name=None)`

Constructs a 1-D convolutional LSTM.

Parameters

- **input_shape** (*Sequence[int]*) – Shape of the inputs excluding batch size.
- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – Sequence of kernel sizes (of length 1), or an int. `kernel_shape` will be expanded to define a kernel size in all dimensions.
- **name** (*Optional[str]*) – Name of the module.

Conv2DLSTM

`class haiku.Conv2DLSTM(input_shape, output_channels, kernel_shape, name=None)`

2-D convolutional LSTM.

The implementation is based on [3]. Given x_t and the previous state (h_{t-1}, c_{t-1}) the core computes

$$\begin{aligned} i_t &= \sigma(W_{ii} * x_t + W_{hi} * h_{t-1} + b_i) \\ f_t &= \sigma(W_{if} * x_t + W_{hf} * h_{t-1} + b_f) \\ g_t &= \tanh(W_{ig} * x_t + W_{hg} * h_{t-1} + b_g) \\ o_t &= \sigma(W_{io} * x_t + W_{ho} * h_{t-1} + b_o) \\ c_t &= f_t c_{t-1} + i_t g_t \\ h_t &= o_t \tanh(c_t) \end{aligned}$$

where $*$ denotes the convolution operator; i_t, f_t, o_t are input, forget and output gate activations, and g_t is a vector of cell updates.

The output is equal to the new hidden state, h_t .

Notes

Forget gate initialization: Following [2] we add 1.0 to b_f after initialization in order to reduce the scale of forgetting in the beginning of the training.

`__init__(input_shape, output_channels, kernel_shape, name=None)`

Constructs a 2-D convolutional LSTM.

Parameters

- **input_shape** (*Sequence[int]*) – Shape of the inputs excluding batch size.
- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – Sequence of kernel sizes (of length 2), or an int. `kernel_shape` will be expanded to define a kernel size in all dimensions.
- **name** (*Optional[str]*) – Name of the module.

Conv3DLSTM

`class haiku.Conv3DLSTM(input_shape, output_channels, kernel_shape, name=None)`

3-D convolutional LSTM.

The implementation is based on [3]. Given x_t and the previous state (h_{t-1}, c_{t-1}) the core computes

$$\begin{aligned} i_t &= \sigma(W_{ii} * x_t + W_{hi} * h_{t-1} + b_i) \\ f_t &= \sigma(W_{if} * x_t + W_{hf} * h_{t-1} + b_f) \\ g_t &= \tanh(W_{ig} * x_t + W_{hg} * h_{t-1} + b_g) \\ o_t &= \sigma(W_{io} * x_t + W_{ho} * h_{t-1} + b_o) \\ c_t &= f_t c_{t-1} + i_t g_t \\ h_t &= o_t \tanh(c_t) \end{aligned}$$

where $*$ denotes the convolution operator; i_t, f_t, o_t are input, forget and output gate activations, and g_t is a vector of cell updates.

The output is equal to the new hidden state, h_t .

Notes

Forget gate initialization: Following [2] we add 1.0 to b_f after initialization in order to reduce the scale of forgetting in the beginning of the training.

`__init__(input_shape, output_channels, kernel_shape, name=None)`

Constructs a 3-D convolutional LSTM.

Parameters

- **input_shape** (*Sequence[int]*) – Shape of the inputs excluding batch size.
- **output_channels** (*int*) – Number of output channels.
- **kernel_shape** (*Union[int, Sequence[int]]*) – Sequence of kernel sizes (of length 3), or an int. `kernel_shape` will be expanded to define a kernel size in all dimensions.
- **name** (*Optional[str]*) – Name of the module.

Attention

MultiHeadAttention

`class haiku.MultiHeadAttention(num_heads, key_size, w_init_scale=None, *, w_init=None, with_bias=True, b_init=None, value_size=None, model_size=None, name=None)`

Multi-headed attention (MHA) module.

This module is intended for attending over sequences of vectors.

Rough sketch: - Compute keys (K), queries (Q), and values (V) as projections of inputs. - Attention weights are computed as $W = \text{softmax}(QK^T / \sqrt{\text{key_size}})$. - Output is another projection of WV^T .

For more detail, see the original Transformer paper: “Attention is all you need” <https://arxiv.org/abs/1706.03762>.

Glossary of shapes: - T: Sequence length. - D: Vector (embedding) size. - H: Number of attention heads.

`__init__`(*num_heads*, *key_size*, *w_init_scale*=None, *, *w_init*=None, *with_bias*=True, *b_init*=None, *value_size*=None, *model_size*=None, *name*=None)

Initialises the module.

Parameters

- **num_heads** (*int*) – Number of independent attention heads (H).
- **key_size** (*int*) – The size of keys (K) and queries used for attention.
- **w_init_scale** (*Optional*[*float*]) – DEPRECATED. Please use *w_init* instead.
- **w_init** (*Optional*[*hk.initializers.Initializer*]) – Initialiser for weights in the linear map. Once *w_init_scale* is fully deprecated *w_init* will become mandatory. Until then it has a default value of *None* for backwards compatability.
- **with_bias** (*bool*) – Whether to add a bias when computing various linear projections.
- **b_init** (*Optional*[*hk.initializers.Initializer*]) – Optional initializer for bias. By default, zero.
- **value_size** (*Optional*[*int*]) – Optional size of the value projection (V). If *None*, defaults to the key size (K).
- **model_size** (*Optional*[*int*]) – Optional size of the output embedding (D'). If *None*, defaults to the key size multiplied by the number of heads (K * H).
- **name** (*Optional*[*str*]) – Optional name for this module.

`__call__`(*query*, *key*, *value*, *mask*=None)

Computes (optionally masked) MHA with queries, keys & values.

This module broadcasts over zero or more ‘batch-like’ leading dimensions.

Parameters

- **query** (*jax.Array*) – Embeddings sequence used to compute queries; shape [..., T', D_q].
- **key** (*jax.Array*) – Embeddings sequence used to compute keys; shape [..., T, D_k].
- **value** (*jax.Array*) – Embeddings sequence used to compute values; shape [..., T, D_v].
- **mask** (*Optional*[*jax.Array*]) – Optional mask applied to attention weights; shape [..., H=1, T', T].

Return type *jax.Array*

Returns

A new sequence of embeddings, consisting of a projection of the attention-weighted value projections; shape [..., T', D'].

Batch

<code>Reshape(output_shape[, preserve_dims, name])</code>	Reshapes input Tensor, preserving the batch dimension.
<code>Flatten([preserve_dims, name])</code>	Flattens the input, preserving the batch dimension(s).
<code>BatchApply(f[, num_dims])</code>	Temporarily merges leading dimensions of input tensors.

Reshape

class `haiku.Reshape(output_shape, preserve_dims=1, name=None)`

Reshapes input Tensor, preserving the batch dimension.

For example, given an input tensor with shape [B, H, W, C, D]:

```
>>> B, H, W, C, D = range(1, 6)
>>> x = jnp.ones([B, H, W, C, D])
```

The default behavior when `output_shape` is `(-1, D)` is to flatten all dimensions between B and D:

```
>>> mod = hk.Reshape(output_shape=(-1, D))
>>> assert mod(x).shape == (B, H*W*C, D)
```

You can change the number of preserved leading dimensions via `preserve_dims`:

```
>>> mod = hk.Reshape(output_shape=(-1, D), preserve_dims=2)
>>> assert mod(x).shape == (B, H, W*C, D)

>>> mod = hk.Reshape(output_shape=(-1, D), preserve_dims=3)
>>> assert mod(x).shape == (B, H, W, C, D)

>>> mod = hk.Reshape(output_shape=(-1, D), preserve_dims=4)
>>> assert mod(x).shape == (B, H, W, C, 1, D)
```

Alternatively, a negative value of `preserve_dims` specifies the number of trailing dimensions to replace with `output_shape`:

```
>>> mod = hk.Reshape(output_shape=(-1, D), preserve_dims=-3)
>>> assert mod(x).shape == (B, H, W*C, D)
```

This is useful in the case of applying the same module to batched and unbatched outputs:

```
>>> mod = hk.Reshape(output_shape=(-1, D), preserve_dims=-3)
>>> assert mod(x[0]).shape == (H, W*C, D)
```

__init__ (`output_shape, preserve_dims=1, name=None`)

Constructs a `Reshape` module.

Parameters

- **output_shape** (`Sequence[int]`) – Shape to reshape the input tensor to while preserving its first `preserve_dims` dimensions. When the special value `-1` appears in `output_shape` the corresponding size is automatically inferred. Note that `-1` can only appear once in `output_shape`. To flatten all non-batch dimensions use `Flatten`.

- **preserve_dims** (*int*) – Number of leading dimensions that will not be reshaped. If negative, this is interpreted instead as the number of trailing dimensions to replace with the new shape.
- **name** (*Optional[str]*) – Name of the module.

Raises ValueError – If `preserve_dims` is zero.

`__call__(inputs)`

Call self as a function.

Flatten

`class haiku.Flatten(preserve_dims=1, name=None)`

Flattens the input, preserving the batch dimension(s).

By default, `Flatten` combines all dimensions except the first. Additional leading dimensions can be preserved by setting `preserve_dims`.

```
>>> x = jnp.ones([3, 2, 4])
>>> flat = hk.Flatten()
>>> flat(x).shape
(3, 8)
```

When the input to flatten has fewer than `preserve_dims` dimensions it is returned unchanged:

```
>>> x = jnp.ones([3])
>>> flat(x).shape
(3,)
```

Alternatively, a negative value of `preserve_dims` specifies the number of trailing dimensions flattened:

```
>>> x = jnp.ones([3, 2, 4])
>>> negative_flat = hk.Flatten(preserve_dims=-2)
>>> negative_flat(x).shape
(3, 8)
```

This allows the same module to be seamlessly applied to a single element or a batch of elements with the same element shape:

```
>> negative_flat(x[0]).shape (8,)
```

`__init__(preserve_dims=1, name=None)`

Constructs a `Reshape` module.

Parameters

- **output_shape** – Shape to reshape the input tensor to while preserving its first `preserve_dims` dimensions. When the special value `-1` appears in `output_shape` the corresponding size is automatically inferred. Note that `-1` can only appear once in `output_shape`. To flatten all non-batch dimensions use `Flatten`.
- **preserve_dims** (*int*) – Number of leading dimensions that will not be reshaped. If negative, this is interpreted instead as the number of trailing dimensions to replace with the new shape.
- **name** (*Optional[str]*) – Name of the module.

Raises ValueError – If `preserve_dims` is zero.

BatchApply

class haiku.**BatchApply**(*f*, *num_dims*=2)

Temporarily merges leading dimensions of input tensors.

Merges the leading dimensions of a tensor into a single dimension, runs the given callable, then splits the leading dimension of the result to match the input.

Input arrays whose rank is smaller than the number of dimensions to collapse are passed unmodified.

This may be useful for applying a module to each timestep of e.g. a [Time, Batch, ...] array.

For some *fs* and platforms, this may be more efficient than `jax.vmap()`, especially when combined with other transformations like `jax.grad()`.

__init__(*f*, *num_dims*=2)

Constructs a *BatchApply* module.

Parameters

- **f** – The callable to be applied to the reshaped array.
- **num_dims** – The number of dimensions to merge.

__call__(*args, **kwargs)

Call self as a function.

Embedding

<code>Embed</code> (<i>vocab_size</i> , <i>embed_dim</i> , ...)	Module for embedding tokens in a low-dimensional space.
<code>EmbedLookupStyle</code> (<i>value</i> [, <i>names</i> , <i>module</i> , ...])	How to return the embedding matrices given IDs.

Embed

class haiku.**Embed**(*vocab_size*=None, *embed_dim*=None, *embedding_matrix*=None, *w_init*=None, *lookup_style*='ARRAY_INDEX', *name*=None, *precision*=<Precision.HIGHEST: 2>)

Module for embedding tokens in a low-dimensional space.

__init__(*vocab_size*=None, *embed_dim*=None, *embedding_matrix*=None, *w_init*=None, *lookup_style*='ARRAY_INDEX', *name*=None, *precision*=<Precision.HIGHEST: 2>)

Constructs an Embed module.

Parameters

- **vocab_size** (*Optional[int]*) – The number of unique tokens to embed. If not provided, an existing vocabulary matrix from which *vocab_size* can be inferred must be provided as *embedding_matrix*.
- **embed_dim** (*Optional[int]*) – Number of dimensions to assign to each embedding. If an existing vocabulary matrix initializes the module, this should not be provided as it will be inferred.
- **embedding_matrix** (*Optional[Union[np.ndarray, jax.Array]]*) – A matrix-like object equivalent in size to [*vocab_size*, *embed_dim*]. If given, it is used as the initial value for the embedding matrix and neither *vocab_size* or *embed_dim* need be

given. If they are given, their values are checked to be consistent with the dimensions of `embedding_matrix`.

- **w_init** (*Optional*[*hk.initializers.Initializer*]) – An initializer for the embeddings matrix. As a default, embeddings are initialized via a truncated normal distribution.
- **lookup_style** (*Union*[*str*, *EmbedLookupStyle*]) – One of the enum values of *EmbedLookupStyle* determining how to access the value of the embeddings given an ID. Regardless the input should be a dense array of integer values representing ids. This setting changes how internally this module maps those ids to embeddings. The result is the same, but the speed and memory tradeoffs are different. It defaults to using NumPy-style array indexing. This value is only the default for the module, and at any given invocation can be overridden in `__call__()`.
- **name** (*Optional*[*str*]) – Optional name for this module.
- **precision** (*jax.lax.Precision*) – Only used when `lookup_style` is `ONE_HOT`. The precision to use for the dot-product between the one-hot-encoded inputs and the embedding vectors. It is possible to attain a ~2x speedup on TPU using *jax.lax.Precision.DEFAULT* at the cost of a slightly lower precision.

Raises ValueError – If none of `embed_dim`, `embedding_matrix` and `vocab_size` are supplied, or if `embedding_matrix` is supplied and `embed_dim` or `vocab_size` is not consistent with the supplied matrix.

`__call__(ids, lookup_style=None, precision=None)`

Lookup embeddings.

Looks up an embedding vector for each value in `ids`. All ids must be within `[0, vocab_size)` to prevent NaNs from propagating.

Parameters

- **ids** (*Union*[*jax.Array*, *Sequence*[*int*]]) – integer array.
- **lookup_style** (*Optional*[*Union*[*str*, *hk.EmbedLookupStyle*]]) – Overrides the `lookup_style` given in the constructor.
- **precision** (*Optional*[*jax.lax.Precision*]) – Overrides the `precision` given in the constructor.

Return type `jax.Array`

Returns Tensor of `ids.shape + [embedding_dim]`.

Raises

- **AttributeError** – If `lookup_style` is not valid.
- **ValueError** – If `ids` is not an integer array.

EmbedLookupStyle

```
class haiku.EmbedLookupStyle(value, names=None, *, module=None, qualname=None, type=None, start=1,
                             boundary=None)
```

How to return the embedding matrices given IDs.

```
ARRAY_INDEX = 1
```

```
ONE_HOT = 2
```

Utilities

<code>Deferred(factory[, call_methods])</code>	Defers the construction of another module until the first call.
--	---

Deferred

```
class haiku.Deferred(factory, call_methods=('__call__',))
```

Defers the construction of another module until the first call.

Deferred can be used to declare modules that depend on computed properties of other modules before those modules are defined. This allows users to separate the declaration and use of modules. For example at the start of your program you can declare two modules which are coupled:

```
>>> encoder = hk.Linear(64)
>>> decoder = hk.Deferred(lambda: hk.Linear(encoder.input_size))
```

Later you can use these naturally (note: that using `decoder` first would cause an error since `encoder.input_size` is only defined after `encoder` has been called):

```
>>> x = jnp.ones([8, 32])
>>> y = encoder(x)
>>> z = decoder(y) # Constructs the Linear encoder by calling the lambda.
```

The result will satisfy the following conditions:

```
>>> assert x.shape == z.shape
>>> assert y.shape == (8, 64)
>>> assert decoder.input_size == encoder.output_size
>>> assert decoder.output_size == encoder.input_size
```

```
__init__(factory, call_methods=('__call__',))
```

Initializes the *Deferred* module.

Parameters

- **factory** (*Callable*[[*T*], *T*]) – A no argument callable which constructs the module to defer to. The first time one of the *call_methods* are called the factory will be run and then the constructed module will be called with the same method and arguments as the deferred module.
- **call_methods** (*Sequence*[*str*]) – Methods which should trigger construction of the target module. The default value configures this module to construct the first time `__call__` is run. If you want to add methods other than call you should explicitly pass them (optionally), for example `call_methods=("__call__", "encode", "decode")`.

property target: T

Returns the target module.

If the factory has not already run this will trigger construction. Subsequent calls to *target* will return the same instance.

Return type T

Returns A *Module* instance as created by the factory function passed into the constructor.

```

__call__(*args, **kwargs)
    Call self as a function.

__setattr__(name, value)
    Implement setattr(self, name, value).

__delattr__(name)
    Implement delattr(self, name).

```

Initializers

<i>Initializer</i>	alias of Callable[[collections.abc.Sequence[int], Any], jax.Array]
<i>Constant</i> (constant)	Initializes with a constant.
<i>Identity</i> ([gain])	Initializer that generates the identity matrix.
<i>Orthogonal</i> ([scale, axis])	Uniform scaling initializer.
<i>RandomNormal</i> ([stddev, mean])	Initializes by sampling from a normal distribution.
<i>RandomUniform</i> ([minval, maxval])	Initializes by sampling from a uniform distribution.
<i>TruncatedNormal</i> ([stddev, mean, lower, upper])	Initializes by sampling from a truncated normal distribution.
<i>VarianceScaling</i> ([scale, mode, distribution, ...])	Initializer which adapts its scale to the shape of the initialized array.
<i>UniformScaling</i> ([scale])	Uniform scaling initializer.

Initializer

haiku.initializers.**Initializer**

alias of Callable[[collections.abc.Sequence[int], Any], jax.Array]

Constant

class haiku.initializers.**Constant**(*constant*)

Initializes with a constant.

```
__init__(constant)
```

Constructs a Constant initializer.

Parameters **constant** (*Union[float, int, complex, np.ndarray, jax.Array]*) – Constant to initialize with.

```
__call__(shape, dtype)
```

Call self as a function.

Return type jax.Array

Identity

class haiku.initializers.**Identity**(*gain=1.0*)

Initializer that generates the identity matrix.

Constructs a 2D identity matrix or batches of these.

__init__(*gain=1.0*)

Constructs an *Identity* initializer.

Parameters *gain* (*Union[float, np.ndarray, jax.Array]*) – Multiplicative factor to apply to the identity matrix.

__call__(*shape, dtype*)

Call self as a function.

Return type *jax.Array*

Orthogonal

class haiku.initializers.**Orthogonal**(*scale=1.0, axis=- 1*)

Uniform scaling initializer.

__init__(*scale=1.0, axis=- 1*)

Construct an initializer for uniformly distributed orthogonal matrices.

These matrices will be row-orthonormal along the access specified by *axis*. If the rank of the weight is greater than 2, the shape will be flattened in all other dimensions and then will be row-orthonormal along the final dimension. Note that this only works if the *axis* dimension is larger, otherwise the matrix will be transposed (equivalently, it will be column orthonormal instead of row orthonormal).

If the shape is not square, the matrices will have orthonormal rows or columns depending on which side is smaller.

Parameters

- **scale** – Scale factor.
- **axis** – Which axis corresponds to the “output dimension” of the tensor.

Returns An orthogonally initialized parameter.

__call__(*shape, dtype*)

Call self as a function.

Return type *jax.Array*

RandomNormal

class haiku.initializers.**RandomNormal**(*stddev=1.0, mean=0.0*)

Initializes by sampling from a normal distribution.

__init__(*stddev=1.0, mean=0.0*)

Constructs a *RandomNormal* initializer.

Parameters

- **stddev** – The standard deviation of the normal distribution to sample from.

- **mean** – The mean of the normal distribution to sample from.

`__call__(shape, dtype)`

Call self as a function.

Return type `jax.Array`

RandomUniform

`class haiku.initializers.RandomUniform(minval=0.0, maxval=1.0)`

Initializes by sampling from a uniform distribution.

`__init__(minval=0.0, maxval=1.0)`

Constructs a *RandomUniform* initializer.

Parameters

- **minval** – The lower limit of the uniform distribution.
- **maxval** – The upper limit of the uniform distribution.

`__call__(shape, dtype)`

Call self as a function.

Return type `jax.Array`

TruncatedNormal

`class haiku.initializers.TruncatedNormal(stddev=1.0, mean=0.0, lower=-2.0, upper=2.0)`

Initializes by sampling from a truncated normal distribution.

`__init__(stddev=1.0, mean=0.0, lower=-2.0, upper=2.0)`

Constructs a *TruncatedNormal* initializer.

Parameters

- **stddev** (*Union[float, jax.Array]*) – The standard deviation parameter of the truncated normal distribution.
- **mean** (*Union[float, complex, jax.Array]*) – The mean of the truncated normal distribution.
- **lower** (*Union[float, jax.Array]*) – Float or array representing the lower bound for truncation.
- **upper** (*Union[float, jax.Array]*) – Float or array representing the upper bound for truncation.

`__call__(shape, dtype)`

Call self as a function.

Return type `jax.Array`

VarianceScaling

```
class haiku.initializers.VarianceScaling(scale=1.0, mode='fan_in', distribution='truncated_normal',
                                         fan_in_axes=None)
```

Initializer which adapts its scale to the shape of the initialized array.

The initializer first computes the scaling factor $s = \text{scale} / n$, where n is:

- Number of input units in the weight tensor, if `mode = fan_in`.
- Number of output units, if `mode = fan_out`.
- Average of the numbers of input and output units, if `mode = fan_avg`.

Then, with `distribution="truncated_normal"` or `"normal"`, samples are drawn from a distribution with a mean of zero and a standard deviation (after truncation, if used) `stddev = sqrt(s)`.

With `distribution=uniform`, samples are drawn from a uniform distribution within `[-limit, limit]`, with `limit = sqrt(3 * s)`.

The variance scaling initializer can be configured to generate other standard initializers using the `scale`, `mode` and `distribution` arguments. Here are some example configurations:

Name	Parameters
<code>glorot_uniform</code>	<code>VarianceScaling(1.0, "fan_avg", "uniform")</code>
<code>glorot_normal</code>	<code>VarianceScaling(1.0, "fan_avg", "truncated_normal")</code>
<code>lecun_uniform</code>	<code>VarianceScaling(1.0, "fan_in", "uniform")</code>
<code>lecun_normal</code>	<code>VarianceScaling(1.0, "fan_in", "truncated_normal")</code>
<code>he_uniform</code>	<code>VarianceScaling(2.0, "fan_in", "uniform")</code>
<code>he_normal</code>	<code>VarianceScaling(2.0, "fan_in", "truncated_normal")</code>

```
__init__(scale=1.0, mode='fan_in', distribution='truncated_normal', fan_in_axes=None)
```

Constructs the `VarianceScaling` initializer.

Parameters

- **scale** – Scale to multiply the variance by.
- **mode** – One of `fan_in`, `fan_out`, `fan_avg`
- **distribution** – Random distribution to use. One of `truncated_normal`, `normal` or `uniform`.
- **fan_in_axes** – Optional sequence of `int` specifying which axes of the shape are part of the fan-in. If none provided, then the weight is assumed to be like a convolution kernel, where all leading dimensions are part of the fan-in, and only the trailing dimension is part of the fan-out. Useful if instantiating multi-headed attention weights.

```
__call__(shape, dtype)
```

Call self as a function.

Return type `jax.Array`

UniformScaling

class haiku.initializers.**UniformScaling**(*scale=1.0*)

Uniform scaling initializer.

Initializes by sampling from a uniform distribution, but with the variance scaled by the inverse square root of the number of input units, multiplied by the scale.

__init__(*scale=1.0*)

Constructs the *UniformScaling* initializer.

Parameters **scale** – Scale to multiply the upper limit of the uniform distribution by.

__call__(*shape, dtype*)

Call self as a function.

Return type jax.Array

Paddings

<i>PadFn</i>	alias of Callable[[int], tuple[int, int]]
<i>is_padfn</i> (padding)	Tests whether the given argument is a single or sequence of PadFns.
<i>create</i> (padding, kernel, rate, n)	Generates the padding required for a given padding algorithm.
<i>create_from_padfn</i> (padding, kernel, rate, n)	Generates the padding required for a given padding algorithm.
<i>create_from_tuple</i> (padding, n)	Create a padding tuple using partially specified padding tuple.
<i>causal</i> (effective_kernel_size)	Pre-padding such that output has no dependence on the future.
<i>full</i> (effective_kernel_size)	Maximal padding whilst not convolving over just padded elements.
<i>reverse_causal</i> (effective_kernel_size)	Post-padding such that output has no dependence on the past.
<i>same</i> (effective_kernel_size)	Pads such that the output size matches input size for stride=1.
<i>valid</i> (effective_kernel_size)	No padding.

PadFn

haiku.pad.**PadFn**

alias of Callable[[int], tuple[int, int]]

is_padfn

`haiku.pad.is_padfn(padding)`

Tests whether the given argument is a single or sequence of PadFns.

Return type bool

create

`haiku.pad.create(padding, kernel, rate, n)`

Generates the padding required for a given padding algorithm.

Parameters

- **padding** (*Union*[*hk.pad.PadFn*, *Sequence*[*hk.pad.PadFn*]]) – callable/tuple or a sequence of callables/tuples. The callables take an integer representing the effective kernel size (kernel size when the rate is 1) and return a sequence of two integers representing the padding before and padding after for that dimension. The tuples are defined with two elements, padding before and after. If *padding* is a sequence it must be of length 1 or *n*.
- **kernel** (*Union*[*int*, *Sequence*[*int*]]) – int or sequence of ints of length *n*. The size of the kernel for each dimension. If it is an int it will be replicated for the non channel and batch dimensions.
- **rate** (*Union*[*int*, *Sequence*[*int*]]) – int or sequence of ints of length *n*. The dilation rate for each dimension. If it is an int it will be replicated for the non channel and batch dimensions.
- **n** (*int*) – the number of spatial dimensions.

Return type *Sequence*[*tuple*[*int*, *int*]]

Returns A sequence of length *n* containing the padding for each element. These are of the form [*pad_before*, *pad_after*].

create_from_padfn

`haiku.pad.create_from_padfn(padding, kernel, rate, n)`

Generates the padding required for a given padding algorithm.

Parameters

- **padding** (*Union*[*hk.pad.PadFn*, *Sequence*[*hk.pad.PadFn*]]) – callable/tuple or a sequence of callables/tuples. The callables take an integer representing the effective kernel size (kernel size when the rate is 1) and return a sequence of two integers representing the padding before and padding after for that dimension. The tuples are defined with two elements, padding before and after. If *padding* is a sequence it must be of length 1 or *n*.
- **kernel** (*Union*[*int*, *Sequence*[*int*]]) – int or sequence of ints of length *n*. The size of the kernel for each dimension. If it is an int it will be replicated for the non channel and batch dimensions.
- **rate** (*Union*[*int*, *Sequence*[*int*]]) – int or sequence of ints of length *n*. The dilation rate for each dimension. If it is an int it will be replicated for the non channel and batch dimensions.
- **n** (*int*) – the number of spatial dimensions.

Return type Sequence[tuple[int, int]]

Returns A sequence of length *n* containing the padding for each element. These are of the form [pad_before, pad_after].

create_from_tuple

haiku.pad.create_from_tuple(padding, n)

Create a padding tuple using partially specified padding tuple.

Return type Sequence[tuple[int, int]]

causal

haiku.pad.causal(effective_kernel_size)

Pre-padding such that output has no dependence on the future.

Return type tuple[int, int]

full

haiku.pad.full(effective_kernel_size)

Maximal padding whilst not convolving over just padded elements.

Return type tuple[int, int]

reverse_causal

haiku.pad.reverse_causal(effective_kernel_size)

Post-padding such that output has no dependence on the past.

Return type tuple[int, int]

same

haiku.pad.same(effective_kernel_size)

Pads such that the output size matches input size for stride=1.

Return type tuple[int, int]

valid

haiku.pad.valid(effective_kernel_size)

No padding.

Return type tuple[int, int]

1.3.5 Full Networks

MLP

```
class haiku.nets.MLP(output_sizes, w_init=None, b_init=None, with_bias=True,
                    activation=<jax._src.custom_derivatives.custom_jvp object>, activate_final=False,
                    name=None)
```

A multi-layer perceptron module.

```
__init__(output_sizes, w_init=None, b_init=None, with_bias=True,
         activation=<jax._src.custom_derivatives.custom_jvp object>, activate_final=False, name=None)
```

Constructs an MLP.

Parameters

- **output_sizes** (*Iterable[int]*) – Sequence of layer sizes.
- **w_init** (*Optional[hk.initializers.Initializer]*) – Initializer for *Linear* weights.
- **b_init** (*Optional[hk.initializers.Initializer]*) – Initializer for *Linear* bias. Must be None if *with_bias*=False.
- **with_bias** (*bool*) – Whether or not to apply a bias in each layer.
- **activation** (*Callable[[jax.Array], jax.Array]*) – Activation function to apply between *Linear* layers. Defaults to ReLU.
- **activate_final** (*bool*) – Whether or not to activate the final layer of the MLP.
- **name** (*Optional[str]*) – Optional name for this module.

Raises ValueError – If *with_bias* is False and *b_init* is not None.

```
__call__(inputs, dropout_rate=None, rng=None)
```

Connects the module to some inputs.

Parameters

- **inputs** (*jax.Array*) – A Tensor of shape [*batch_size*, *input_size*].
- **dropout_rate** (*Optional[float]*) – Optional dropout rate.
- **rng** – Optional RNG key. Require when using dropout.

Return type *jax.Array*

Returns The output of the model of size [*batch_size*, *output_size*].

```
reverse(activate_final=None, name=None)
```

Returns a new MLP which is the layer-wise reverse of this MLP.

NOTE: Since computing the reverse of an MLP requires knowing the input size of each linear layer this method will fail if the module has not been called at least once.

The contract of reverse is that the reversed module will accept the output of the parent module as input and produce an output which is the input size of the parent.

```
>>> mlp = hk.nets.MLP([1, 2, 3])
>>> mlp_in = jnp.ones([1, 2])
>>> y = mlp(mlp_in)
>>> rev = mlp.reverse()
```

(continues on next page)

(continued from previous page)

```
>>> rev_mlp_out = rev(y)
>>> mlp_in.shape == rev_mlp_out.shape
True
```

Parameters

- **activate_final** (*Optional[bool]*) – Whether the final layer of the MLP should be activated.
- **name** (*Optional[str]*) – Optional name for the new module. The default name will be the name of the current module prefixed with "reversed_".

Return type *MLP*

Returns An MLP instance which is the reverse of the current instance. Note these instances do not share weights and, apart from being symmetric to each other, are not coupled in any way.

MobileNet**MobileNetV1**

```
class haiku.nets.MobileNetV1(strides=(1, 2, 1, 2, 1, 2, 1, 1, 1, 1, 1, 2, 1), channels=(64, 128, 128, 256, 256,
512, 512, 512, 512, 512, 512, 1024, 1024), num_classes=1000, use_bn=True,
name=None)
```

MobileNetV1 model.

```
__init__(strides=(1, 2, 1, 2, 1, 2, 1, 1, 1, 1, 1, 2, 1), channels=(64, 128, 128, 256, 256, 512, 512, 512, 512,
512, 512, 1024, 1024), num_classes=1000, use_bn=True, name=None)
```

Constructs a MobileNetV1 model.

Parameters

- **strides** (*Sequence[int]*) – The stride to use the in depthwise convolution in each mobilenet block.
- **channels** (*Sequence[int]*) – Number of output channels from the pointwise convolution to use in each block.
- **num_classes** (*int*) – Number of classes.
- **use_bn** (*bool*) – Whether or not to use batch normalization. Defaults to True. When true, biases are not used. When false, biases are used.
- **name** (*Optional[str]*) – Name of the module.

```
__call__(inputs, is_training)
```

Call self as a function.

Return type `jax.Array`

ResNet

<code>ResNet(blocks_per_group, num_classes[, ...])</code>	ResNet model.
<code>ResNet.BlockGroup(channels, num_blocks, ...)</code>	Higher level block for ResNet implementation.
<code>ResNet.BlockV1(channels, stride, ...[, name])</code>	ResNet V1 block with optional bottleneck.
<code>ResNet.BlockV2(channels, stride, ...[, name])</code>	ResNet V2 block with optional bottleneck.
<code>ResNet18(num_classes[, bn_config, ...])</code>	ResNet18.
<code>ResNet34(num_classes[, bn_config, ...])</code>	ResNet34.
<code>ResNet50(num_classes[, bn_config, ...])</code>	ResNet50.
<code>ResNet101(num_classes[, bn_config, ...])</code>	ResNet101.
<code>ResNet152(num_classes[, bn_config, ...])</code>	ResNet152.
<code>ResNet200(num_classes[, bn_config, ...])</code>	ResNet200.

ResNet

```
class haiku.nets.ResNet(blocks_per_group, num_classes, bn_config=None, resnet_v2=False,
    bottleneck=True, channels_per_group=(256, 512, 1024, 2048),
    use_projection=(True, True, True, True), logits_config=None, name=None,
    initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet model.

```
class BlockGroup(channels, num_blocks, stride, bn_config, resnet_v2, bottleneck, use_projection,
    name=None)
```

Higher level block for ResNet implementation.

```
__call__(inputs, is_training, test_local_stats)
```

Call self as a function.

```
__init__(channels, num_blocks, stride, bn_config, resnet_v2, bottleneck, use_projection, name=None)
```

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters **name** (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If **name** is not provided then the class name for the current instance is converted to **lower_snake_case** and used instead.

```
class BlockV1(channels, stride, use_projection, bn_config, bottleneck, name=None)
```

ResNet V1 block with optional bottleneck.

```
__call__(inputs, is_training, test_local_stats)
```

Call self as a function.

```
__init__(channels, stride, use_projection, bn_config, bottleneck, name=None)
```

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters **name** (*Optional[str]*) – An optional string name for the class. Must be a valid Python identifier. If **name** is not provided then the class name for the current instance is converted to **lower_snake_case** and used instead.

```
class BlockV2(channels, stride, use_projection, bn_config, bottleneck, name=None)
```

ResNet V2 block with optional bottleneck.

`__call__(inputs, is_training, test_local_stats)`

Call self as a function.

`__init__(channels, stride, use_projection, bn_config, bottleneck, name=None)`

Initializes the current module with the given name.

Subclasses should call this constructor before creating other modules or variables such that those modules are named correctly.

Parameters `name` (*Optional*[*str*]) – An optional string name for the class. Must be a valid Python identifier. If `name` is not provided then the class name for the current instance is converted to `lower_snake_case` and used instead.

`__init__(blocks_per_group, num_classes, bn_config=None, resnet_v2=False, bottleneck=True, channels_per_group=(256, 512, 1024, 2048), use_projection=(True, True, True, True), logits_config=None, name=None, initial_conv_config=None, strides=(1, 2, 2, 2))`

Constructs a ResNet model.

Parameters

- **blocks_per_group** (*Sequence*[*int*]) – A sequence of length 4 that indicates the number of blocks created in each group.
- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional*[*Mapping*[*str*, *FloatStrOrBool*]]) – A dictionary of two elements, `decay_rate` and `eps` to be passed on to the [BatchNorm](#) layers. By default the `decay_rate` is 0.9 and `eps` is 1e-5.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to `False`.
- **bottleneck** (*bool*) – Whether the block should bottleneck or not. Defaults to `True`.
- **channels_per_group** (*Sequence*[*int*]) – A sequence of length 4 that indicates the number of channels used for each block in each group.
- **use_projection** (*Sequence*[*bool*]) – A sequence of length 4 that indicates whether each residual block should use projection.
- **logits_config** (*Optional*[*Mapping*[*str*, *Any*]]) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional*[*str*]) – Name of the module.
- **initial_conv_config** (*Optional*[*Mapping*[*str*, *FloatStrOrBool*]]) – Keyword arguments passed to the constructor of the initial [Conv2D](#) module.
- **strides** (*Sequence*[*int*]) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

`__call__(inputs, is_training, test_local_stats=False)`

Call self as a function.

ResNet18

```
class haiku.nets.ResNet18(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                          name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet18.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
         initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, `decay_rate` and `eps` to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to `False`.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional[str]*) – Name of the module.
- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial *Conv2D* module.
- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

ResNet34

```
class haiku.nets.ResNet34(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                          name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet34.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
         initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, `decay_rate` and `eps` to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to `False`.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional[str]*) – Name of the module.
- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial *Conv2D* module.

- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

ResNet50

```
class haiku.nets.ResNet50(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                          name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet50.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
         initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, `decay_rate` and `eps` to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to `False`.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional[str]*) – Name of the module.
- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial *Conv2D* module.
- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

ResNet101

```
class haiku.nets.ResNet101(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                            name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet101.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
         initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, `decay_rate` and `eps` to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to `False`.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional[str]*) – Name of the module.

- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial *Conv2D* module.
- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

ResNet152

```
class haiku.nets.ResNet152(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                           name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet152.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
          initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, *decay_rate* and *eps* to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to *False*.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.
- **name** (*Optional[str]*) – Name of the module.
- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial *Conv2D* module.
- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

ResNet200

```
class haiku.nets.ResNet200(num_classes, bn_config=None, resnet_v2=False, logits_config=None,
                           name=None, initial_conv_config=None, strides=(1, 2, 2, 2))
```

ResNet200.

```
__init__(num_classes, bn_config=None, resnet_v2=False, logits_config=None, name=None,
          initial_conv_config=None, strides=(1, 2, 2, 2))
```

Constructs a ResNet model.

Parameters

- **num_classes** (*int*) – The number of classes to classify the inputs into.
- **bn_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – A dictionary of two elements, *decay_rate* and *eps* to be passed on to the *BatchNorm* layers.
- **resnet_v2** (*bool*) – Whether to use the v1 or v2 ResNet implementation. Defaults to *False*.
- **logits_config** (*Optional[Mapping[str, Any]]*) – A dictionary of keyword arguments for the logits layer.

- **name** (*Optional[str]*) – Name of the module.
- **initial_conv_config** (*Optional[Mapping[str, FloatStrOrBool]]*) – Keyword arguments passed to the constructor of the initial `Conv2D` module.
- **strides** (*Sequence[int]*) – A sequence of length 4 that indicates the size of stride of convolutions for each block in each group.

VectorQuantizer

<code>VectorQuantizer(embedding_dim, ..., dtype, ...)</code>	Haiku module representing the VQ-VAE layer.
<code>VectorQuantizerEMA(embedding_dim, ..., dtype, ...)</code>	Haiku module representing the VQ-VAE layer.

VectorQuantizer

```
class haiku.nets.VectorQuantizer(embedding_dim, num_embeddings, commitment_cost, dtype=<class
                                'jax.numpy.float32'>, name=None, cross_replica_axis=None)
```

Haiku module representing the VQ-VAE layer.

Implements the algorithm presented in “Neural Discrete Representation Learning” by van den Oord et al. <https://arxiv.org/abs/1711.00937>

Input any tensor to be quantized. Last dimension will be used as space in which to quantize. All other dimensions will be flattened and will be seen as different examples to quantize.

The output tensor will have the same shape as the input.

For example a tensor with shape [16, 32, 32, 64] will be reshaped into [16384, 64] and all 16384 vectors (each of 64 dimensions) will be quantized independently.

embedding_dim

integer representing the dimensionality of the tensors in the quantized space. Inputs to the modules must be in this format as well.

num_embeddings

integer, the number of vectors in the quantized space.

commitment_cost

scalar which controls the weighting of the loss terms (see equation 4 in the paper - this variable is Beta).

```
__init__(embedding_dim, num_embeddings, commitment_cost, dtype=<class 'jax.numpy.float32'>,
          name=None, cross_replica_axis=None)
```

Initializes a VQ-VAE module.

Parameters

- **embedding_dim** (*int*) – dimensionality of the tensors in the quantized space. Inputs to the modules must be in this format as well.
- **num_embeddings** (*int*) – number of vectors in the quantized space.
- **commitment_cost** (*float*) – scalar which controls the weighting of the loss terms (see equation 4 in the paper - this variable is Beta).
- **dtype** (*Any*) – dtype for the embeddings variable, defaults to `float32`.
- **name** (*Optional[str]*) – name of the module.

- **cross_replica_axis** (*Optional[str]*) – If not `None`, it should be a string representing the axis name over which this module is being run within a `jax.pmap()`. Supplying this argument means that perplexity is calculated across all replicas on that axis.

`__call__` (*inputs, is_training*)

Connects the module to some inputs.

Parameters

- **inputs** – Tensor, final dimension must be equal to `embedding_dim`. All other leading dimensions will be flattened and treated as a large batch.
- **is_training** – boolean, whether this connection is to training data.

Returns

Dictionary containing the following keys and values:

- **quantize**: Tensor containing the quantized version of the input.
- **loss**: Tensor containing the loss to optimize.
- **perplexity**: Tensor containing the perplexity of the encodings.
- **encodings**: Tensor containing the discrete encodings, ie which element of the quantized space each input element was mapped to.
- **encoding_indices**: Tensor containing the discrete encoding indices, ie which element of the quantized space each input element was mapped to.

Return type dict

quantize (*encoding_indices*)

Returns embedding tensor for a batch of indices.

VectorQuantizerEMA

```
class haiku.nets.VectorQuantizerEMA(embedding_dim, num_embeddings, commitment_cost, decay,
                                   epsilon=1e-05, dtype=<class 'jax.numpy.float32'>,
                                   cross_replica_axis=None, name=None)
```

Haiku module representing the VQ-VAE layer.

Implements a slightly modified version of the algorithm presented in “Neural Discrete Representation Learning” by van den Oord et al. <https://arxiv.org/abs/1711.00937>

The difference between *VectorQuantizerEMA* and *VectorQuantizer* is that this module uses *ExponentialMovingAverages* to update the embedding vectors instead of an auxiliary loss. This has the advantage that the embedding updates are independent of the choice of optimizer (SGD, RMSProp, Adam, K-Fac, ...) used for the encoder, decoder and other parts of the architecture. For most experiments the EMA version trains faster than the non-EMA version.

Input any tensor to be quantized. Last dimension will be used as space in which to quantize. All other dimensions will be flattened and will be seen as different examples to quantize.

The output tensor will have the same shape as the input.

For example a tensor with shape [16, 32, 32, 64] will be reshaped into [16384, 64] and all 16384 vectors (each of 64 dimensions) will be quantized independently.

embedding_dim

integer representing the dimensionality of the tensors in the quantized space. Inputs to the modules must be in this format as well.

num_embeddings

integer, the number of vectors in the quantized space.

commitment_cost

scalar which controls the weighting of the loss terms (see equation 4 in the paper).

decay

float, decay for the moving averages.

epsilon

small float constant to avoid numerical instability.

```
__init__(embedding_dim, num_embeddings, commitment_cost, decay, epsilon=1e-05, dtype=<class 'jax.numpy.float32'>, cross_replica_axis=None, name=None)
```

Initializes a VQ-VAE EMA module.

Parameters

- **embedding_dim** – integer representing the dimensionality of the tensors in the quantized space. Inputs to the modules must be in this format as well.
- **num_embeddings** – integer, the number of vectors in the quantized space.
- **commitment_cost** – scalar which controls the weighting of the loss terms (see equation 4 in the paper - this variable is Beta).
- **decay** – float between 0 and 1, controls the speed of the Exponential Moving Averages.
- **epsilon** (*float*) – small constant to aid numerical stability, default 1e-5.
- **dtype** (*Any*) – dtype for the embeddings variable, defaults to `float32`.
- **cross_replica_axis** (*Optional[str]*) – If not `None`, it should be a string representing the axis name over which this module is being run within a `jax.pmap()`. Supplying this argument means that cluster statistics and the perplexity are calculated across all replicas on that axis.
- **name** (*Optional[str]*) – name of the module.

```
__call__(inputs, is_training)
```

Connects the module to some inputs.

Parameters

- **inputs** – Tensor, final dimension must be equal to `embedding_dim`. All other leading dimensions will be flattened and treated as a large batch.
- **is_training** – boolean, whether this connection is to training data. When this is set to `False`, the internal moving average statistics will not be updated.

Returns

Dictionary containing the following keys and values:

- **quantize**: Tensor containing the quantized version of the input.
- **loss**: Tensor containing the loss to optimize.
- **perplexity**: Tensor containing the perplexity of the encodings.

- **encodings**: Tensor containing the discrete encodings, ie which element of the quantized space each input element was mapped to.
- **encoding_indices**: Tensor containing the discrete encoding indices, ie which element of the quantized space each input element was mapped to.

Return type dict

quantize(*encoding_indices*)

Returns embedding tensor for a batch of indices.

1.3.6 JAX Fundamentals

Control Flow

<code>cond(pred, true_fun, false_fun, *operands[, ...])</code>	Conditionally apply <code>true_fun</code> or <code>false_fun</code> .
<code>fori_loop(lower, upper, body_fun, init_val)</code>	Equivalent to <code>jax.lax.fori_loop()</code> with Haiku state passed in/out.
<code>map(f, xs)</code>	Equivalent to <code>jax.lax.map()</code> but with Haiku state passed in/out.
<code>scan(f, init, xs[, length, reverse, unroll])</code>	Equivalent to <code>jax.lax.scan()</code> but with Haiku state passed in/out.
<code>switch(index, branches, *operands)</code>	Equivalent to <code>jax.lax.switch()</code> but with Haiku state passed in/out.
<code>while_loop(cond_fun, body_fun, init_val)</code>	Equivalent to <code>jax.lax.while_loop</code> with Haiku state threaded in/out.

cond

`haiku.cond(pred, true_fun, false_fun, *operands, operand=<object object>, linear=None)`

Conditionally apply `true_fun` or `false_fun`.

Wraps XLA's [Conditional](#) operator.

Provided arguments are correctly typed, `cond()` has equivalent semantics to this Python implementation, where `pred` must be a scalar type:

```
def cond(pred, true_fun, false_fun, *operands):
    if pred:
        return true_fun(*operands)
    else:
        return false_fun(*operands)
```

In contrast with `jax.lax.select()`, using `cond` indicates that only one of the two branches is executed (up to compiler rewrites and optimizations). However, when transformed with `vmap()` to operate over a batch of predicates, `cond` is converted to `select()`.

Parameters

- **pred** – Boolean scalar type, indicating which branch function to apply.
- **true_fun** (Callable) – Function (A -> B), to be applied if `pred` is True.
- **false_fun** (Callable) – Function (A -> B), to be applied if `pred` is False.

- **operands** – Operands (A) input to either branch depending on `pred`. The type can be a scalar, array, or any pytree (nested Python tuple/list/dict) thereof.

Returns Value (B) of either `true_fun(*operands)` or `false_fun(*operands)`, depending on the value of `pred`. The type can be a scalar, array, or any pytree (nested Python tuple/list/dict) thereof.

fori_loop

`haiku.fori_loop(lower, upper, body_fun, init_val)`

Equivalent to `jax.lax.fori_loop()` with Haiku state passed in/out.

map

`haiku.map(f, xs)`

Equivalent to `jax.lax.map()` but with Haiku state passed in/out.

scan

`haiku.scan(f, init, xs, length=None, reverse=False, unroll=1)`

Equivalent to `jax.lax.scan()` but with Haiku state passed in/out.

switch

`haiku.switch(index, branches, *operands)`

Equivalent to `jax.lax.switch()` but with Haiku state passed in/out.

Note that creating parameters inside a switch branch is not supported, as such at init time we recommend you unconditionally evaluate all branches of your switch and only use the switch at apply. For example:

```
>>> experts = [hk.nets.MLP([300, 100, 10]) for _ in range(5)]
>>> x = jnp.ones([1, 28 * 28])
>>> if hk.running_init():
...     # During init unconditionally create params/state for all experts.
...     for expert in experts:
...         out = expert(x)
... else:
...     # During apply conditionally apply (and update) only one expert.
...     index = jax.random.randint(hk.next_rng_key(), [], 0, len(experts) - 1)
...     out = hk.switch(index, experts, x)
```

Parameters

- **index** – Integer scalar type, indicating which branch function to apply.
- **branches** – Sequence of functions (A -> B) to be applied based on index.
- **operands** – Operands (A) input to whichever branch is applied.

Returns Value (B) of `branch(*operands)` for the branch that was selected based on index.

while_loop

`haiku.while_loop(cond_fun, body_fun, init_val)`

Equivalent to `jax.lax.while_loop` with Haiku state threaded in/out.

JAX Transforms

<code>eval_shape(fun, *args, **kwargs)</code>	Equivalent to <code>jax.eval_shape</code> with any changed Haiku state discarded.
<code>grad(fun[, argnums, has_aux, holomorphic])</code>	Creates a function which evaluates the gradient of <code>fun</code> .
<code>remat(fun, *[, prevent_cse, policy, ...])</code>	Equivalent to <code>jax.checkpoint</code> but passing Haiku state.
<code>value_and_grad(fun[, argnums, has_aux, ...])</code>	Creates a function which evaluates both <code>fun</code> and the grad of <code>fun</code> .
<code>vmap(fun[, in_axes, out_axes, axis_name, ...])</code>	Equivalent to <code>jax.vmap()</code> with module parameters/state not mapped.

eval_shape

`haiku.eval_shape(fun, *args, **kwargs)`

Equivalent to `jax.eval_shape` with any changed Haiku state discarded.

grad

`haiku.grad(fun, argnums=0, has_aux=False, holomorphic=False)`

Creates a function which evaluates the gradient of `fun`.

NOTE: You only need this in a very specific case that you want to take a gradient **inside** a `transform()`ed function and the function you are differentiating uses `set_state()`. For example:

```
>>> class MyModule(hk.Module):
...     def __call__(self, x):
...         hk.set_state("last", x ** 2)
...         return x ** 2
```

```
>>> def f(x):
...     m = MyModule()
...     g = hk.grad(m)(x)
...     return g
```

```
>>> f = hk.transform_with_state(f)
>>> x = jnp.array(2.)
>>> params, state = jax.jit(f.init)(None, x)
>>> print(state["my_module"]["last"])
4.0
```

Parameters

- **fun** – Function to be differentiated. Its arguments at positions specified by `argnums` should be arrays, scalars, or standard Python containers. It should return a scalar (which includes arrays with shape `()` but not arrays with shape `(1,)` etc.)

- **argnums** – Optional, integer or tuple of integers. Specifies which positional argument(s) to differentiate with respect to (default 0).
- **has_aux** – Optional, bool. Indicates whether `fun` returns a pair where the first element is considered the output of the mathematical function to be differentiated and the second element is auxiliary data. Default False.
- **holomorphic** – Optional, bool. Indicates whether `fun` is promised to be holomorphic. Default False.

Returns A function with the same arguments as `fun`, that evaluates the gradient of `fun`. If `argnums` is an integer then the gradient has the same shape and type as the positional argument indicated by that integer. If `argnums` is a tuple of integers, the gradient is a tuple of values with the same shapes and types as the corresponding arguments. If `has_aux` is True then a pair of `gradient`, `auxiliary_data` is returned.

For example:

```
>>> grad_tanh = jax.grad(jax.numpy.tanh)
>>> print(grad_tanh(0.2))
0.96...
```

remat

`haiku.remat(fun, *, prevent_cse=True, policy=None, static_argnums=())`

Equivalent to `jax.checkpoint` but passing Haiku state.

Return type Callable

value_and_grad

`haiku.value_and_grad(fun, argnums=0, has_aux=False, holomorphic=False)`

Creates a function which evaluates both `fun` and the grad of `fun`.

NOTE: You only need this in a very specific case that you want to take a gradient **inside** a `transform()`d function and the function you are differentiating uses `set_state()`. For example:

```
>>> class MyModule(hk.Module):
...     def __call__(self, x):
...         hk.set_state("last", jnp.sum(x))
...         return x ** 2
```

```
>>> def f(x):
...     m = MyModule()
...     y, g = hk.value_and_grad(m)(x)
...     return y, g
```

```
>>> f = hk.transform_with_state(f)
>>> x = jnp.array(2.)
>>> _ = jax.jit(f.init)(None, x)
```

Parameters

- **fun** – Function to be differentiated. Its arguments at positions specified by `argnums` should be arrays, scalars, or standard Python containers. It should return a scalar (which includes arrays with shape `()` but not arrays with shape `(1,)` etc.)
- **argnums** – Optional, integer or tuple of integers. Specifies which positional argument(s) to differentiate with respect to (default 0).
- **has_aux** – Optional, bool. Indicates whether `fun` returns a pair where the first element is considered the output of the mathematical function to be differentiated and the second element is auxiliary data. Default False.
- **holomorphic** – Optional, bool. Indicates whether `fun` is promised to be holomorphic. Default False.

Returns A function with the same arguments as `fun` that evaluates both `fun` and the gradient of `fun` and returns them as a pair (a two-element tuple). If `argnums` is an integer then the gradient has the same shape and type as the positional argument indicated by that integer. If `argnums` is a tuple of integers, the gradient is a tuple of values with the same shapes and types as the corresponding arguments.

vmap

`haiku.vmap(fun, in_axes=0, out_axes=0, axis_name=None, axis_size=None, *, split_rng)`

Equivalent to `jax.vmap()` with module parameters/state not mapped.

The behaviour of Haiku random key APIs under `vmap()` is controlled by the `split_rng` argument:

```
>>> x = jnp.arange(2)
>>> f = hk.vmap(lambda _: hk.next_rng_key(), split_rng=False)
>>> key1, key2 = f(x)
>>> assert (key1 == key2).all()
```

```
>>> f = hk.vmap(lambda _: hk.next_rng_key(), split_rng=True)
>>> key1, key2 = f(x)
>>> assert not (key1 == key2).all()
```

Random numbers in Haiku are typically used for two things, firstly for initialising model parameters, and secondly for creating random samples as part of the forward pass of a neural network (e.g. for dropout). If you are using `vmap()` with a module that uses Haiku random keys for both (e.g. you don't pass keys explicitly into the network), then it is quite likely that you will want to vary the value of `split_rng` depending on whether we are initializing (e.g. creating model parameters) or applying the model. An easy way to do this is to set `split_rng=(not hk.running_init())`.

For more advanced use cases, such as mapping module parameters, we suggest users instead use `lift()` or `transparent_lift()` in combination with `jax.vmap()`.

Parameters

- **fun** (*Callable[... Any]*) – See `jax.vmap()`.
- **in_axes** – See `jax.vmap()`.
- **out_axes** – See `jax.vmap()`.
- **axis_name** (*Optional[str]*) – See `jax.vmap()`.
- **axis_size** (*Optional[int]*) – See `jax.vmap()`.

- **split_rng** (*bool*) – Controls whether random key APIs in Haiku (e.g. `next_rng_key()`) return different (aka. the internal key is split before calling your mapped function) or the same (aka. the internal key is broadcast before calling your mapped function) key. See the docstring for examples.

Return type Callable[..., Any]

Returns See `jax.vmap()`.

1.3.7 Mixed Precision

Automatic Mixed Precision

<code>set_policy(cls, policy)</code>	Uses the given policy for all instances of the module class.
<code>current_policy()</code>	Retrieves the currently active policy in the current context.
<code>get_policy(cls)</code>	Retrieves the currently active policy for the given class.
<code>clear_policy(cls)</code>	Clears any policy associated with the given class.
<code>push_policy(cls, policy)</code>	Sets the given policy for the given class while the context is active.

set_policy

`haiku.mixed_precision.set_policy(cls, policy)`

Uses the given policy for all instances of the module class.

NOTE: Policies are only applied to modules created in the current thread.

A mixed precision policy describes how inputs, module parameters and module outputs should be cast at runtime. By applying a policy to a given type of module, you can control how all instances of that module behave in your program.

For example, you might want to try running a ResNet50 model in a mixture of `float16` and `float32` on GPU to get higher throughput. To do so you can apply a mixed precision policy to the ResNet50 type that will create parameters in `float32`, but cast them to `float16` before use, along with all module inputs:

```
>>> policy = jmp.get_policy('params=float32,compute=float16,output=float32')
>>> hk.mixed_precision.set_policy(hk.nets.ResNet50, policy)
>>> net = hk.nets.ResNet50(4)
>>> x = jnp.ones([4, 224, 224, 3])
>>> print(net(x, is_training=True))
[[0. 0. 0. 0.]
 [0. 0. 0. 0.]
 [0. 0. 0. 0.]
 [0. 0. 0. 0.]
```

For a fully worked mixed precision example see the imagenet example in Haiku's examples directory. This example shows mixed precision on GPU offering a 2x speedup in training time with only a small impact on final top-1 accuracy.

```
>>> hk.mixed_precision.clear_policy(hk.nets.ResNet50)
```

Parameters

- **cls** (*type[hk.Module]*) – A Haiku module class.
- **policy** (*jmp.Policy*) – A JMP policy to apply to the module.

See also:

- *push_policy()*: Context manager for setting policies.
- *current_policy()*: Retrieves the currently active policy (if any).
- *clear_policy()*: Clears any policies associated with a class.
- *get_policy()*: Gets the policy for a given class.

current_policy

`haiku.mixed_precision.current_policy()`

Retrieves the currently active policy in the current context.

Return type `Optional[jmp.Policy]`

Returns The currently active mixed precision policy, or `None`.

See also:

- *clear_policy()*: Clears any policies associated with a class.
- *get_policy()*: Gets the policy for a given class.
- *set_policy()*: Sets a policy for a given class.
- *push_policy()*: Context manager for setting policies.

get_policy

`haiku.mixed_precision.get_policy(cls)`

Retrieves the currently active policy for the given class.

Note that policies applied explicitly to a top level class (e.g. `ResNet`) will be applied implicitly to all child modules (e.g. `ConvND`) called from the parent. This function only returns policies that have been applied explicitly (e.g. via *set_policy()*).

Parameters **cls** (*type[hk.Module]*) – A Haiku module class.

Return type `Optional[jmp.Policy]`

Returns A JMP policy that is used for the given class, or `None` if one is not active.

See also:

- *current_policy()*: Retrieves the currently active policy (if any).
- *clear_policy()*: Clears any policies associated with a class.
- *set_policy()*: Sets a policy for a given class.
- *push_policy()*: Context manager for setting policies.

clear_policy

`haiku.mixed_precision.clear_policy(cls)`

Clears any policy associated with the given class.

Parameters `cls` (*type*[`hk.Module`]) – A Haiku module class.

See also:

- `current_policy()`: Retrieves the currently active policy (if any).
- `get_policy()`: Gets the policy for a given class.
- `set_policy()`: Sets a policy for a given class.
- `push_policy()`: Context manager for setting policies.

push_policy

`haiku.mixed_precision.push_policy(cls, policy)`

Sets the given policy for the given class while the context is active.

Parameters

- `cls` (*type*[`hk.Module`]) – A Haiku module class.
- `policy` (*jmp.Policy*) – A JMP policy to apply to the module.

Yields None.

See also:

- `clear_policy()`: Clears any policies associated with a class.
- `get_policy()`: Gets the policy for a given class.
- `set_policy()`: Sets a policy for a given class.
- `current_policy()`: Retrieves the currently active policy (if any).

1.3.8 Experimental

Graphviz Visualisation

<code>abstract_to_dot(fun)</code>	Converts a function using Haiku modules to a dot graph.
-----------------------------------	---

abstract_to_dot

`haiku.experimental.abstract_to_dot(fun)`

Converts a function using Haiku modules to a dot graph.

Same as `to_dot()` but uses JAX's abstract interpretation machinery to evaluate the function without requiring concrete inputs. Valid inputs for the wrapped function include `jax.ShapeDtypeStruct`.

`abstract_to_dot()` does not support data-dependent control-flow, because no concrete values are provided to the function.

(continued from previous page)

```

| mlp/~linear_2 (Linear) | Linear(output_size=10, name='linear_2') | w: f32[100,
↪10] | f32[8,100] | f32[8,10] | 1,010 | 4.04 KB |
|   ^ mlp (MLP) | | | | | b: f32[10] ↪
↪ | | | | |
+-----+-----+-----+-----+-----+-----+
↪ - - - - - + - - - - - + - - - - - + - - - - - + - - - - - +

```

Possible values for columns:

- **module**: Displays module and method name.
- **config**: Displays the constructor arguments used for the module.
- **owned_params**: Displays parameters directly owned by this module.
- **input**: Displays module inputs.
- **output**: Displays module output.
- **params_size**: Displays the number of parameters
- **params_bytes**: Displays parameter size in bytes.

Possible values for filters:

- **has_output**: Only include methods returning a value other than None.
- **has_params**: Removes methods from modules that do not have parameters.

Parameters

- **f** (*Union[Callable[... , Any], hk.Transformed, hk.TransformedWithState]*) – A function to transform OR one of the init/apply functions from Haiku or the result of `transform()` or `transform_with_state()`.
- **columns** (*Optional[Sequence[str]]*) – A list of column names to enable.
- **filters** (*Optional[Sequence[str]]*) – A list of filters to apply to remove certain module methods.
- **tabulate_kwargs** – Keyword arguments to pass to `tabulate.tabulate(...)`.

Return type Callable[... , str]**Returns** A callable that takes the same arguments as `f` but returns a string summarising the modules used during the execution of `f`.**See also:**`eval_summary()`: Raw data used to generate this table.

eval_summary

`haiku.experimental.eval_summary(f)`

Records module method calls performed by `f`.

```

>>> f = lambda x: hk.nets.MLP([300, 100, 10])(x)
>>> x = jnp.ones([8, 28 * 28])
>>> for i in hk.experimental.eval_summary(f)(x):
...   print("mod := {:14} | in := {} out := {}".format(
...       i.module_details.module.module_name, i.args_spec[0], i.output_spec))
mod := mlp                | in := f32[8,784] out := f32[8,10]
mod := mlp/~linear_0     | in := f32[8,784] out := f32[8,300]
mod := mlp/~linear_1     | in := f32[8,300] out := f32[8,100]
mod := mlp/~linear_2     | in := f32[8,100] out := f32[8,10]

```

Parameters `f` (*Union[Callable[... Any], hk.Transformed, hk.TransformedWithState]*) – A function or transformed function to trace.

Return type `Callable[... Sequence[MethodInvocation]]`

Returns A callable taking the same arguments as the provided function, but returning a sequence of *MethodInvocation* instances revealing the methods called on each module when applying `f`.

See also:

`tabulate()`: Pretty prints a summary of the execution of a function.

ArraySpec

`class haiku.experimental.ArraySpec(shape, dtype)`

Shaped and sized specification of an array.

shape

Shape of the array.

Type `Sequence[int]`

dtype

DType of the array.

Type `jnp.dtype`

`__delattr__(name)`

Implement `delattr(self, name)`.

`__eq__(other)`

Return `self==value`.

`__hash__()`

Return `hash(self)`.

`__init__(shape, dtype)`

`__setattr__(name, value)`

Implement `setattr(self, name, value)`.

MethodInvocation

class haiku.experimental.**MethodInvocation**(*module_details*, *args_spec*, *kwargs_spec*, *output_spec*, *context*, *call_stack*)

Record of a method being invoked on a given module.

module_details

Details about which module and method were invoked.

Type *ModuleDetails*

args_spec

Positional arguments to the method invocation with arrays replaced by *ArraySpec*.

Type tuple[Any, ...]

kwargs_spec

Keyword arguments to the method invocation with arrays replaced by *ArraySpec*.

Type dict[str, Any]

output_spec

Output of the method invocation with arrays replaced by *ArraySpec*.

Type Any

context

Additional context information for the method call as provided by `intercept_methods()`.

Type `hk.MethodContext`

call_stack

Stack of modules currently active while calling this module method. For example if A calls B which calls C then the call stack for C will be [B_DETAILS, A_DETAILS].

Type Sequence[*ModuleDetails*]

__delattr__(*name*)

Implement `delattr(self, name)`.

__eq__(*other*)

Return `self==value`.

__hash__()

Return `hash(self)`.

__init__(*module_details*, *args_spec*, *kwargs_spec*, *output_spec*, *context*, *call_stack*)

__setattr__(*name*, *value*)

Implement `setattr(self, name, value)`.

ModuleDetails

`class haiku.experimental.ModuleDetails(module, method_name, params, state)`

Module and method related information.

module

A *Module* instance.

Type `hk.Module`

method_name

The method name that was invoked on the module.

Type `str`

params

The modules params dict with arrays converted to *ArraySpec*.

Type `Mapping[str, ArraySpec]`

state

The modules state dict with arrays converted to *ArraySpec*.

Type `Mapping[str, ArraySpec]`

__delattr__(name)

Implement `delattr(self, name)`.

__eq__(other)

Return `self==value`.

__hash__()

Return `hash(self)`.

__init__(module, method_name, params, state)

__setattr__(name, value)

Implement `setattr(self, name, value)`.

Managing State

`check_jax_usage([enabled])`

Ensures JAX APIs (e.g.

check_jax_usage

`haiku.experimental.check_jax_usage(enabled=True)`

Ensures JAX APIs (e.g. `jax.vmap()`) are used correctly with Haiku.

JAX transforms (like `jax.vmap()`) and control flow (e.g. `jax.lax.cond()`) expect pure functions to be passed in. Some functions in Haiku (for example `get_parameter()`) have side effects and thus functions using them are only pure after using `transform()` (et al).

Sometimes it is convenient to use JAX transforms or control flow before transforming your function (for example, to `vmap()` the application of a module) but when doing so you need to be careful to use the Haiku overloaded version of the underlying JAX function, which carefully makes the function(s) you pass in pure functions before calling the underlying JAX function.

`check_jax_usage()` enables checking raw JAX transforms are used appropriately inside Haiku transformed functions. Incorrect usage of JAX transforms will result in an error.

Consider the function below, it is not a pure function (a function of its inputs with no side effects) because we call into a Haiku API (`get_parameter()`) which during init will create a parameter and register it with Haiku.

```
>>> def f():
...     return hk.get_parameter("some_param", [], init=jnp.zeros)
```

We should not use this with JAX APIs like `jax.vmap()` (because it is not a pure function). `check_jax_usage()` allows you to tell Haiku to make incorrect usages of JAX APIs an error:

```
>>> previous_value = hk.experimental.check_jax_usage(True)
>>> jax.vmap(f, axis_size=2)()
Traceback (most recent call last):
...
haiku.JaxUsageError: ...
```

Using the Haiku wrapped version works correctly:

```
>>> print(hk.vmap(f, axis_size=2, split_rng=False)())
[0. 0.]
```

Parameters `enabled` (*bool*) – Boolean indicating whether usage should be checked or not.

Return type `bool`

Returns Boolean with the previous value for this setting.

Optimizations

<code>optimize_rng_use(fun)</code>	Optimizes a RNG key splitting in <code>fun</code> .
<code>module_auto_repr(enabled)</code>	Disables automatically generating an implementation of <code>Module.__repr__</code> .
<code>fast_eval_shape(fun, *args, **kwargs)</code>	Equivalent to <code>eval_shape</code> in JAX.
<code>rng_reserve_size(size)</code>	Change amount of RNG keys reserved when calling <code>next_rng_key</code> .

`optimize_rng_use`

`haiku.experimental.optimize_rng_use(fun)`

Optimizes a RNG key splitting in `fun`.

Our strategy here is to use abstract interpretation to run your function twice, the first time we use `jax.eval_shape()` to avoid spending any flops and simply observe how many times you call `next_rng_key()`. We then run your function again, but this time we reserve enough RNG keys ahead of time such that we only need to call `jax.random.split()` once.

In the following example, we need three random samples for our weight matrices in our 3-layer MLP. To draw these samples we use `next_rng_key()` which will split a new key for each sample. By using `optimize_rng_use()` Haiku will pre-allocate exactly enough RNGs for `f` to be evaluated by splitting the input key once and only once. For large models (unlike this example) this can lead to a reduction in compilation time of both `init` and `apply`, with `init` seeing a larger expected speedup as it performs more RNG key splitting in general.

```
>>> def f(x):
...     net = hk.nets.MLP([300, 100, 10])
...     return net(x)
>>> f = hk.experimental.optimize_rng_use(f)
>>> f = hk.transform(f)
>>> params = f.init(jax.random.PRNGKey(42), jnp.ones([1, 1]))
```

Parameters `fun` – A function to wrap.

Returns A function that applies `fun` but only requires one call to `jax.random.split()` by Haiku.

module_auto_repr

`haiku.experimental.module_auto_repr(enabled)`

Disables automatically generating an implementation of `Module.__repr__`.

By default, Haiku will automatically generate a useful string representation of modules for printing. For example:

```
>>> print(hk.Linear(1))
Linear(output_size=1)
```

In some cases, objects passed into module constructors may be slow to print, for example very nested data structures, or you may be rapidly creating and throwing away modules (e.g. in a test) and don't want to pay the overhead of converting to string.

This config option enables users to disable the automatic repr feature globally in Haiku:

```
>>> previous_value = hk.experimental.module_auto_repr(False)
>>> print(hk.Linear(1))
<...Linear object at ...>
```

```
>>> previous_value = hk.experimental.module_auto_repr(True)
>>> print(hk.Linear(1))
Linear(output_size=1)
```

To disable the feature on a per-subclass basis assign `AUTO_REPR = False` as a property on your class, for example:

```
>>> class NoAutoRepr(hk.Module):
...     AUTO_REPR = False
>>> print(NoAutoRepr())
<...NoAutoRepr object at ...>
```

Parameters `enabled` (*bool*) – Boolean indicating whether a module should be enabled.

Return type `bool`

Returns The previous value of this config setting.

fast_eval_shape

`haiku.experimental.fast_eval_shape(fun, *args, **kwargs)`

Equivalent to `eval_shape` in JAX.

This utility is equivalent to `eval_shape` in JAX except that it avoids running Haiku functions whose shapes are trivially known. This can avoid some Python overheads in JAX which can accumulate for very large models.

Optimizations:

- All parameter/state initialisers replaced with zeros.
- `hk.dropout` replaced with identity.
- `jax.random.fold_in` replaced with identity.

Parameters

- **fun** – The function to trace.
- ***args** – Positional arguments to `fun`.
- ****kwargs** – Keyword arguments to `fun`.

Returns The shape produced by `fun` for the given args/kwargs.

rng_reserve_size

`haiku.experimental.rng_reserve_size(size)`

Change amount of RNG keys reserved when calling `next_rng_key`.

Parameters `size` (*int*) – amount of keys to reserve when splitting off a key through `next_rng_key()`, defaults to 1. Reserving larger blocks of keys can improve compilation and run-time of your model. Changing the reservation size will change RNG keys returned by `next_rng_key`, and will change the generated random numbers.

Return type `int`

Returns The previous value of the `rng_reserve_size` setting.

jaxpr_info

<code>make_model_info(f[, name, ...])</code>	Creates a function that computes flop, param and state information.
<code>as_html(module[, min_flop, outvars, last])</code>	Formats a <i>Module</i> as a tree of interactive HTML elements.
<code>as_html_page(module[, min_flop])</code>	Formats the output of <code>make_model_info</code> as an interactive HTML page.
<code>css()</code>	The CSS for HTML visualization of a <i>Module</i> .
<code>format_module(module[, depth])</code>	Recursively formats module information as a human readable string.
<code>js()</code>	The JavaScript for HTML visualization of a <i>Module</i> .
<code>Expression(primitive, invars, outvars[, ...])</code>	Information about a single JAX expression.
<code>Module(name[, flops, expressions, ...])</code>	Information about a Haiku module.

make_model_info

```
haiku.experimental.jaxpr_info.make_model_info(f, name=None, include_module_info=True,
                                             compute_flops=None, axis_env=None)
```

Creates a function that computes flop, param and state information.

Parameters

- **f** (*Callable*[... , *Any*]) – The function for which to compute information. Haiku modules and `jax.named_call` expressions will be represented as nested `Module`s in the result.
- **name** (*Optional*[*str*]) – Optional, the name of the root expression.
- **include_module_info** (*bool*) – Whether to include parameter and state count information for haiku modules. Can be slow for very large computations.
- **compute_flops** (*Optional*[*ComputeFlopsFn*]) – Optional, a function that returns an estimate of the number of flops required to execute an equation.
- **axis_env** (*Optional*[*Sequence*[*tuple*[*Any*, *int*]]]) – Sizes of pmapped axes. See docs of `jax.make_jaxpr` for details.

Return type `Callable`[... , `Module`]

Returns A wrapped version of `f` that when applied to example arguments returns a `Module` representation of `f` for those arguments.

`Module` and `Expression` contain high level information about JAX operations (jaxprs) and can be visualized in concise and interactive formats; see `format_module`, `as_html_page` or `as_html`.

as_html

```
haiku.experimental.jaxpr_info.as_html(module, min_flop=1000, outvars="", last=False)
```

Formats a `Module` as a tree of interactive HTML elements.

When embedding this in a page, the outputs of `css` and `js` must be embedded too for the visualization to work. To only visualize a single module directly, see `as_html_page`.

Parameters

- **module** (`Module`) – The module to visualize, as an interactive HTML tree.
- **min_flop** (*int*) – Minimum number of flops for an operation to be shown.
- **outvars** (*str*) – For internal use, the outputs of this module.
- **last** (*bool*) – For internal use, whether this module is the last of its siblings.

Return type `str`

Returns HTML representation of `module`.

as_html_page

`haiku.experimental.jaxpr_info.as_html_page(module, min_flop=1000)`

Formats the output of `make_model_info` as an interactive HTML page.

Return type str

css

`haiku.experimental.jaxpr_info.css()`

The CSS for HTML visualization of a *Module*.

Return type str

format_module

`haiku.experimental.jaxpr_info.format_module(module, depth=0)`

Recursively formats module information as a human readable string.

Return type str

js

`haiku.experimental.jaxpr_info.js()`

The JavaScript for HTML visualization of a *Module*.

Return type str

Expression

`class haiku.experimental.jaxpr_info.Expression(primitive, invars, outvars, flops=None, details="", params=<factory>, submodule=None, first_outvar="", name_stack=<factory>)`

Information about a single JAX expression.

Module

`class haiku.experimental.jaxpr_info.Module(name, flops=None, expressions=<factory>, total_param_size=0, param_info=<factory>, total_state_size=0, state_info=<factory>)`

Information about a Haiku module.

1.3.9 Configuration

<code>context(*[, check_jax_usage, ...])</code>	Context manager for setting config options.
<code>set(*[, check_jax_usage, module_auto_repr, ...])</code>	Sets the given config option(s).

context

```
haiku.config.context(*, check_jax_usage=None, module_auto_repr=None, restore_flatmap=None,
                    rng_reserve_size=None)
```

Context manager for setting config options.

This context manager can be used to override config settings in a given context, values that are not explicitly passed as keyword arguments retain their current value:

```
>>> with hk.config.context(check_jax_usage=True):
...     pass
```

Parameters

- **check_jax_usage** (*Optional[bool]*) – Checks that jax transforms and control flow are used appropriately in Haiku transformed functions.
- **module_auto_repr** (*Optional[bool]*) – Can be used to disable the “to string” functionality that is part of Haiku’s base constructor.
- **restore_flatmap** (*Optional[bool]*) – Whether legacy checkpoints should be restored in the old FlatMap datatype (as returned by `to_immutable_dict`), default is to restore these as plain dicts.
- **rng_reserve_size** (*Optional[int]*) – amount of keys to reserve when splitting off a key through `next_rng_key()`, defaults to 1. Reserving larger blocks of keys can improve compilation and run-time of your model. Changing the reservation size will change RNG keys returned by `next_rng_key`, and will change the generated random numbers.

Returns Context manager that applies the given configs while active.

set

```
haiku.config.set(*, check_jax_usage=None, module_auto_repr=None, restore_flatmap=None,
                rng_reserve_size=None)
```

Sets the given config option(s).

```
>>> hk.config.set(module_auto_repr=False)
>>> hk.Linear(1)
<...Linear object at ...>
>>> hk.config.set(module_auto_repr=True)
>>> hk.Linear(1)
Linear(output_size=1)
```

Parameters

- **check_jax_usage** (*Optional[bool]*) – Checks that jax transforms and control flow are used appropriately in Haiku transformed functions.

- **module_auto_repr** (*Optional[bool]*) – Can be used to disable the “to string” functionality that is part of Haiku’s base constructor.
- **restore_flatmap** (*Optional[bool]*) – Whether legacy checkpoints should be restored in the old FlatMap datatype (as returned by `to_immutable_dict`), default is to restore these as plain dicts.
- **rng_reserve_size** (*Optional[int]*) – amount of keys to reserve when splitting off a key through `next_rng_key()`, defaults to 1. Reserving larger blocks of keys can improve compilation and run-time of your model. Changing the reservation size will change RNG keys returned by `next_rng_key`, and will change the generated random numbers.

1.3.10 Utilities

Data Structures

<code>filter</code> (predicate, structure)	Filters an input structure according to a user specified predicate.
<code>is_subset</code> (*, subset, superset)	Checks whether the leaves of subset appear in superset.
<code>map</code> (fn, structure)	Maps a function to an input structure accordingly.
<code>merge</code> (*structures[, check_duplicates])	Merges multiple input structures.
<code>partition</code> (predicate, structure)	Partitions the input structure in two according to a given predicate.
<code>partition_n</code> (fn, structure, n)	Partitions a structure into <i>n</i> structures.
<code>to_haiku_dict</code> (structure)	Returns a copy of the given two level structure.
<code>to_immutable_dict</code> (mapping)	Returns an immutable copy of the given mapping.
<code>to_mutable_dict</code> (mapping)	Turns an immutable FlatMapping into a mutable dict.
<code>traverse</code> (structure)	Iterates over a structure yielding module names, names and values.
<code>tree_bytes</code> (tree)	Sums the size in bytes of all arrays in a pytree.
<code>tree_size</code> (tree)	Sums the sizes of all arrays in a pytree.

filter

`haiku.data_structures.filter`(predicate, structure)

Filters an input structure according to a user specified predicate.

```
>>> params = {'linear': {'w': None, 'b': None}}
>>> predicate = lambda module_name, name, value: name == 'w'
>>> hk.data_structures.filter(predicate, params)
{'linear': {'w': None}}
```

Note: returns a new structure not a view.

Parameters

- **predicate** (*Callable[[str, str, T], bool]*) – criterion to be used to partition the input data. The `predicate` argument is expected to be a boolean function taking as inputs the name of the module, the name of a given entry in the module data bundle (e.g. parameter name) and the corresponding data.
- **structure** (*Mapping[str, Mapping[str, T]]*) – Haiku params or state data structure to be filtered.

Return type Mapping[str, Mapping[str, T]]

Returns All the input parameters or state as selected by the input predicate.

is_subset

haiku.data_structures.is_subset(*, subset, superset)

Checks whether the leaves of subset appear in superset.

Note that this is vacuously true in the case that both structures have no leaves:

```
>>> hk.data_structures.is_subset(subset={'a': {}}, superset={})
True
```

Parameters

- **subset** (*Mapping[str, Mapping[str, Any]]*) – The subset to check.
- **superset** (*Mapping[str, Mapping[str, Any]]*) – The superset to check.

Return type bool

Returns A boolean indicating whether all elements in subset are contained in superset.

map

haiku.data_structures.map(fn, structure)

Maps a function to an input structure accordingly.

```
>>> params = {'linear': {'w': 1.0, 'b': 2.0}}
>>> fn = lambda module_name, name, value: 2 * value if name == 'w' else value
>>> hk.data_structures.map(fn, params)
{'linear': {'b': 2.0, 'w': 2.0}}
```

Note: returns a new structure not a view.

Parameters

- **fn** (*Callable[[str, str, InT], OutT]*) – criterion to be used to map the input data. The *fn* argument is expected to be a function taking as inputs the name of the module, the name of a given entry in the module data bundle (e.g. parameter name) and the corresponding data, and returning a new value.
- **structure** (*Mapping[str, Mapping[str, InT]]*) – Haiku params or state data structure to be mapped.

Return type Mapping[str, Mapping[str, OutT]]

Returns All the input parameters or state as mapped by the input *fn*.

merge

`haiku.data_structures.merge(*structures, check_duplicates=False)`

Merges multiple input structures.

```
>>> weights = {'linear': {'w': None}}
>>> biases = {'linear': {'b': None}}
>>> hk.data_structures.merge(weights, biases)
{'linear': {'w': None, 'b': None}}
```

When structures are not disjoint the output will contain the value from the last structure for each path:

```
>>> weights1 = {'linear': {'w': 1}}
>>> weights2 = {'linear': {'w': 2}}
>>> hk.data_structures.merge(weights1, weights2)
{'linear': {'w': 2}}
```

Note: returns a new structure not a view.

Parameters

- ***structures** – One or more structures to merge.
- **check_duplicates** (*bool*) – If True, a `ValueError` will be thrown if an array is found in multiple structures but with a different shape and dtype.

Return type `MutableMapping[str, MutableMapping[str, Any]]`

Returns A single structure with an entry for each path in the input structures.

partition

`haiku.data_structures.partition(predicate, structure)`

Partitions the input structure in two according to a given predicate.

For a given set of parameters, you can use `partition()` to split them:

```
>>> params = {'linear': {'w': None, 'b': None}}
>>> predicate = lambda module_name, name, value: name == 'w'
>>> weights, biases = hk.data_structures.partition(predicate, params)
>>> weights
{'linear': {'w': None}}
>>> biases
{'linear': {'b': None}}
```

Note: returns new structures not a view.

Parameters

- **predicate** (`Callable[[str, str, jax.Array], bool]`) – criterion to be used to partition the input data. The `predicate` argument is expected to be a boolean function taking as inputs the name of the module, the name of a given entry in the module data bundle (e.g. parameter name) and the corresponding data.
- **structure** (`Mapping[str, Mapping[str, T]]`) – Haiku params or state data structure to be partitioned.

Return type `tuple[Mapping[str, Mapping[str, T]], Mapping[str, Mapping[str, T]]]`

Returns

A tuple containing all the params or state as partitioned by the input predicate. Entries matching the predicate will be in the first structure, and the rest will be in the second.

partition_n

`haiku.data_structures.partition_n(fn, structure, n)`

Partitions a structure into n structures.

For a given set of parameters, you can use `partition_n()` to split them into n groups. For example, to split your parameters/gradients by module name:

```
>>> def partition_by_module(structure):
...     cnt = itertools.count()
...     d = collections.defaultdict(lambda: next(cnt))
...     fn = lambda m, n, v: d[m]
...     return hk.data_structures.partition_n(fn, structure, len(structure))
```

```
>>> structure = {'layer_{i}': {'w': None, 'b': None} for i in range(3)}
>>> for substructure in partition_by_module(structure):
...     print(substructure)
{'layer_0': {'b': None, 'w': None}}
{'layer_1': {'b': None, 'w': None}}
{'layer_2': {'b': None, 'w': None}}
```

Parameters

- **fn** (*Callable*[[*str*, *str*, *T*], *int*]) – Callable returning which bucket in $[0, n)$ the given element should be output.
- **structure** (*Mapping*[*str*, *Mapping*[*str*, *T*]]) – Haiku params or state data structure to be partitioned.
- **n** (*int*) – The total number of buckets.

Return type `tuple[Mapping[str, Mapping[str, T]], ...]`

Returns A tuple of size n , where each element will contain the values for which the function returned the current index.

to_haiku_dict

`haiku.data_structures.to_haiku_dict(structure)`

Returns a copy of the given two level structure.

Uses the same mapping type as Haiku will return from `init` or `apply` functions.

Parameters **structure** (*Mapping*[*K*, *V*]) – A two level mapping to copy.

Return type `MutableMapping[K, V]`

Returns A new two level mapping with the same contents as the input.

to_immutable_dict

`haiku.data_structures.to_immutable_dict(mapping)`

Returns an immutable copy of the given mapping.

Return type Mapping[K, V]

to_mutable_dict

`haiku.data_structures.to_mutable_dict(mapping)`

Turns an immutable FlatMapping into a mutable dict.

traverse

`haiku.data_structures.traverse(structure)`

Iterates over a structure yielding module names, names and values.

NOTE: Items are iterated in key sorted order.

Parameters `structure` (*Mapping[str, Mapping[str, T]]*) – The structure to traverse.

Yields Tuples of the module name, name and value from the given structure.

Return type Generator[tuple[str, str, T], None, None]

tree_bytes

`haiku.data_structures.tree_bytes(tree)`

Sums the size in bytes of all arrays in a pytree.

Note that this is the minimum size of the array (e.g. for a float32 we need at least 4 bytes) however on some accelerators buffers may occupy more memory due to padding/alignment constraints.

For example given a ResNet50 model:

```

>>> f = hk.transform_with_state(lambda x: hk.nets.ResNet50(1000)(x, True))
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([128, 224, 224, 3])
>>> params, state = f.init(rng, x)

```

We can count the number of parameters and their size at f32:

```

>>> num_params = hk.data_structures.tree_size(params)
>>> byte_size = hk.data_structures.tree_bytes(params)
>>> print(f'{num_params} params, size: {byte_size / 1e6:.2f}MB')
25557032 params, size: 102.23MB

```

And compare that with casting our parameters to bf16:

```

>>> params = jax.tree_util.tree_map(lambda x: x.astype(jnp.bfloat16), params)
>>> num_params = hk.data_structures.tree_size(params)
>>> byte_size = hk.data_structures.tree_bytes(params)
>>> print(f'{num_params} params, size: {byte_size / 1e6:.2f}MB')
25557032 params, size: 51.11MB

```

Parameters `tree` – A tree of `jax.Arrays`.

Return type `int`

Returns The total size in bytes of the array(s) in the input.

`tree_size`

`haiku.data_structures.tree_size(tree)`

Sums the sizes of all arrays in a pytree.

For example given a ResNet50 model:

```
>>> f = hk.transform_with_state(lambda x: hk.nets.ResNet50(1000)(x, True))
>>> rng = jax.random.PRNGKey(42)
>>> x = jnp.ones([128, 224, 224, 3])
>>> params, state = f.init(rng, x)
```

We can count the number of parameters and their size at f32:

```
>>> num_params = hk.data_structures.tree_size(params)
>>> byte_size = hk.data_structures.tree_bytes(params)
>>> print(f'{num_params} params, size: {byte_size / 1e6:.2f}MB')
25557032 params, size: 102.23MB
```

And compare that with casting our parameters to bf16:

```
>>> params = jax.tree_util.tree_map(lambda x: x.astype(jnp.bfloat16), params)
>>> num_params = hk.data_structures.tree_size(params)
>>> byte_size = hk.data_structures.tree_bytes(params)
>>> print(f'{num_params} params, size: {byte_size / 1e6:.2f}MB')
25557032 params, size: 51.11MB
```

Parameters `tree` – A tree of `jax.Arrays`.

Return type `int`

Returns The total size (number of elements) of the array(s) in the input.

Testing

<code>transform_and_run([f, seed, run_apply, ...])</code>	Transforms the given function and runs init then (optionally) apply.
---	--

transform_and_run

```
haiku.testing.transform_and_run(f=None, seed=42, run_apply=True, jax_transform=None, *,
                               map_rng=None)
```

Transforms the given function and runs init then (optionally) apply.

Equivalent to:

```
>>> def f(x):
...     return x
>>> x = jnp.ones([])
>>> rng = jax.random.PRNGKey(42)
>>> f = hk.transform_with_state(f)
>>> params, state = f.init(rng, x)
>>> out = f.apply(params, state, rng, x)
```

This function makes it very convenient to unit test Haiku:

```
>>> class MyTest(unittest.TestCase):
...     @hk.testing.transform_and_run
...     def test_linear_output(self):
...         mod = hk.Linear(1)
...         out = mod(jnp.ones([1, 1]))
...         self.assertEqual(out.ndim, 2)
```

It can also be combined with `chex` to test all pure/jit/pmap versions of a function:

```
>>> class MyTest(unittest.TestCase):
...     @chex.all_variants
...     def test_linear_output(self):
...         @hk.testing.transform_and_run(jax_transform=self.variant)
...         def f(inputs):
...             mod = hk.Linear(1)
...             return mod(inputs)
...         out = f(jnp.ones([1, 1]))
...         self.assertEqual(out.ndim, 2)
```

And can also be useful in an interactive environment like `ipython`, `Jupyter` or `Google Colaboratory`:

```
>>> f = lambda x: hk.Bias()(x)
>>> print(hk.testing.transform_and_run(f)(jnp.ones([1, 1])))
[[1.]]
```

See `transform()` for more details.

To use this with `pmap` (without `chex`) you need to additionally pass in a function to map the init/apply rng keys. For example, if you want every instance of your `pmap` to have the same key:

```
>>> def same_key_on_all_devices(key):
...     return jnp.broadcast_to(key, (jax.local_device_count(), *key.shape))
```

```
>>> @hk.testing.transform_and_run(jax_transform=jax.pmap,
...                               map_rng=same_key_on_all_devices)
... def test_something():
...     ...
```

Or you can use a different key:

```
>>> def different_key_on_all_devices(key):
...     return jax.random.split(key, jax.local_device_count())

>>> @hk.testing.transform_and_run(jax_transform=jax.pmap,
...                               map_rng=different_key_on_all_devices)
... def test_something_else():
...     ...
```

Parameters

- **f** (*Optional*[*Fn*]) – A function method to transform.
- **seed** (*Optional*[*int*]) – A seed to pass to `init` and `apply`.
- **run_apply** (*bool*) – Whether to run `apply` as well as `init`. Defaults to `true`.
- **jax_transform** (*Optional*[*Callable*[[*Fn*], *Fn*]]) – An optional jax transform to apply on the `init` and `apply` functions.
- **map_rng** (*Optional*[*Callable*[[*Key*], *Key*]]) – If set to a non-None value broadcast the `init/apply` rngs broadcast_rng-ways.

Return type T

Returns A function that `transform()`s `f` and runs `init` and optionally `apply`.

Conditional Computation

<code>running_init()</code>	Return True if running the <code>init</code> function of a Haiku transform.
-----------------------------	---

running_init

`haiku.running_init()`

Return True if running the `init` function of a Haiku transform.

In general you should not need to gate behaviour of your module based on whether you are running `init` or `apply`, but sometimes (e.g. when making use of JAX control flow) this is required.

For example, if you want to use `switch()` to pick between experts, when we run your `init` function we need to ensure that params/state for all experts are created (unconditionally) but during `apply` we want to conditionally `apply` (and perhaps update the internal state) of only one of our experts:

```
>>> experts = [hk.nets.ResNet50(10) for _ in range(5)]
>>> x = jnp.ones([1, 224, 224, 3])
>>> if hk.running_init():
...     # During init unconditionally create params/state for all experts.
...     for expert in experts:
...         out = expert(x, is_training=True)
... else:
...     # During apply conditionally apply (and update) only one expert.
...     index = jax.random.randint(hk.next_rng_key(), [], 0, len(experts) - 1)
...     out = hk.switch(index, experts, x)
```

Return type bool

Returns True if running `init` otherwise False.

Functions

<code><i>multinomial</i></code> (rng, logits, num_samples)	Draws samples from a multinomial distribution.
<code><i>one_hot</i></code> (x, num_classes[, dtype])	Returns a one-hot version of indices.

multinomial

`haiku.multinomial`(rng, logits, num_samples)

Draws samples from a multinomial distribution.

DEPRECATED: Use `jax.random.categorical` instead.

Parameters

- **rng** – A JAX PRNGKey.
- **logits** – Unnormalized log-probabilities, where last dimension is categories.
- **num_samples** – Number of samples to draw.

Returns Chosen categories, of shape `logits.shape[:-1] + (num_samples,)`.

one_hot

`haiku.one_hot`(x, num_classes, dtype=<class 'jax.numpy.float32'>)

Returns a one-hot version of indices.

DEPRECATED: Use `jax.nn.one_hot(x, num_classes).astype(dtype)` instead.

Parameters

- **x** – A tensor of indices.
- **num_classes** – Number of classes in the one-hot dimension.
- **dtype** – The dtype.

Returns

The one-hot tensor. If indices' shape is `[A, B, ...]`, shape is `[A, B, ... num_classes]`.

1.3.11 References

1.4 Haiku and Flax interop

Utilities to move seamlessly between Haiku and Flax.

1.4.1 Flax inside Haiku

Using a Flax module inside a `hk.transform` (or `hk.transform_with_state`) is straight forward.

First construct an instance of your module, and then use `hkflax.lift` to “lift” (see `[hk.lift]`) the parameters and state from the Flax module into the Haiku transform.

Example:

```
[ ]: import jax
import jax.numpy as jnp
import haiku as hk
import haiku.experimental.flax as hkflax
import flax.linen as flax_nn

def f(x):
    mod = hkflax.lift(flax_nn.Dense(10), name='my_flax_module')
    x = mod(x)
    return x

f = hk.transform(f)
x = jnp.ones([1, 1])
rng = jax.random.PRNGKey(42)
params = f.init(rng, x) # params contains the parameters for MyFlaxModule.
f.apply(params, None, x) # MyFlaxModule will be passed parameters from params.

No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and rerun for more
↳info.)

Array([[ 0.33030465, -1.3496182,  0.02847686, -1.6579462, -0.9166192,
         0.2883583, -0.046898,  0.6414894, -0.404975, -2.1162813 ]],
↳dtype=float32)
```

To use a stateful module simply swap `hk.transform` for `hk.transform_with_state`.

1.4.2 Haiku inside Flax

There are two supported approaches for converting Haiku code to Flax. Both produce a Flax linen `nn.Module` which encapsulates the Haiku code and provides `init` and `apply` methods to create and use parameters and state.

- Convert an `hk.Module` to `nn.Module` `<#hk-Module>`__``.
- Convert an `hk.transform` to `nn.Module` `<#hk-transform>`__``.
- Convert an `hk.transform_with_state` to `nn.Module` `<#hk-transform>`__``.

Converting `hk.Module`

For stateless modules you simply need to construct the Flax module via `hkflax.Module.create`:

```
[ ]: mod = hkflax.Module.create(hk.Linear, 1) # hk.Linear(1)
```

You can use this like a regular Flax `nn.Module` (because it is one!):

```
[ ]: rng = jax.random.PRNGKey(42)
      x = jnp.ones([1, 1])
      variables = mod.init(rng, x)
      out = mod.apply(variables, x)
```

For a stateful module like ResNet, you need to also handle output state, again this is the same as stateful Flax modules:

```
[ ]: mod = hkflax.Module.create(hk.nets.ResNet50, 10)

# Regular Flax code from here on:
rng = jax.random.PRNGKey(42)
x = jnp.ones([1, 224, 224, 3])
variables = mod.init(rng, x, is_training=True)
for _ in range(10):
    out, state_out = mod.apply(variables, x, is_training=True,
                               mutable=['state'])
    variables = {**variables, **state_out}
```

Converting `hk.transform` or `hk.transform_with_state`

`hkflax.Module` can be created from the result of `hk.transform` or `hk.transform_with_state` if you prefer:

```
[ ]: def mlp(x):
      x = hk.Linear(300)(x)
      x = hk.Linear(100)(x)
      x = hk.Linear(10)(x)
      return x

mlp = hk.transform(mlp)
mlp = hkflax.Module(mlp)

rng = jax.random.PRNGKey(42)
x = jnp.ones([1, 28 * 28])
variables = mlp.init(rng, x)
out = mlp.apply(variables, x)
```

Gotchas

Initialization is different

Flax and Haiku take different approaches to RNG key splitting. As such at the moment the parameters returned from `hkflax.Module(f).init` will differ from `hk.transform(f).init`.

We have a route to support making Haiku transform match initialization of the Flax module, but there is not a path for the opposite direction at the moment.

If aligning initialization across Haiku and Flax is important to you, we recommend using one of the libraries to create parameters, and then manipulate the params/state dictionary to match the other library as needed:

```
# Utilities.
import haiku.data_structures as hkds
```

(continues on next page)

```

make_flat = {f'{m}/{n}': w for m, n, w in hkds.traverse(d)}

def make_nested(d):
    out = {}
    for k, w in d.items():
        m, n = k.rsplit('/', 1)
        out.setdefault(m, {})
        out[m][n] = w
    return out

# The two modules here should be equivalent when run with Flax or Haiku.
f = hk.transform_with_state(...)
flax_mod = hkflax.Module(f)

# Option 1: Convert Haiku initialized params/state to Flax.
params, state = f.init(...)
variables = {'params': make_flat(params), 'state': make_flat(state)}

# Option 2: Convert Flax initialized variables to Haiku.
variables = flax_mod.init(...)
params = make_nested(variables.get('params', {}))
state = make_nested(variables.get('state', {}))

# The output of the Haiku transformed function or the Flax function should be
# equivalent with either init.
out, state = f.apply(params, state, ...)
out, variables_out = flax_mod.apply(variables, ..., mutable=['state'])

```

Multiple forward methods

hkflax.Module only support `__call__` at the moment, please let us know if this is a blocker for you.

1.5 Haiku and jax2tf

jax2tf is an advanced JAX feature supporting staging JAX programs out as TensorFlow graphs.

This is a useful feature if you want to integrate with an existing TensorFlow codebase or tool. In this tutorial we will demonstrate defining a simple model in Haiku, converting it to TensorFlow as a `tf.Module` and then training it.

We'll then save the model as a TensorFlow `SavedModel` so it can be used later in other TensorFlow programs.

```
[1]: !pip install dm-tree dm-sonnet tensorflow tensorflow_datasets ipywidgets matplotlib >/
    ↪ dev/null
```

```
[2]: import haiku as hk
import jax
import jax.numpy as jnp
from jax.experimental import jax2tf
import sonnet as snt
```

(continues on next page)

(continued from previous page)

```
import tensorflow as tf
import tree
```

1.5.1 Define your model in JAX

First things first, we need to define our model using Haiku and JAX. For MNIST we can use a trivial model like an MLP.

We initialize the model using JAX and get initial parameter values. If you wanted you could additionally go on to train your model using JAX, but in this example we will do that in TensorFlow.

```
[3]: def f(x):
      net = hk.nets.MLP([300, 100, 10])
      return net(x)
```

```
f = hk.transform(f)
```

```
rng = jax.random.PRNGKey(42)
x = jnp.ones([1, 28 * 28 * 1])
params = f.init(rng, x)
```

```
WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and
↳rerun for more info.)
```

1.5.2 Convert to TensorFlow

TensorFlow ships with a module abstraction that supports common tasks like collecting model parameters.

Sonnet is a library of `tf.Module` subclasses including common NN layers, optimizers and some metrics. Sonnet is a sister library to Haiku developed by the same team.

We will use Sonnet's module class for some nice `name_scope`-ing and later we will use the Adam optimizer implemented in Sonnet as well as some utility functions.

```
[4]: def create_variable(path, value):
      name = '/'.join(map(str, path)).replace('~', '_')
      return tf.Variable(value, name=name)

class JaxModule(snt.Module):
    def __init__(self, params, apply_fn, name=None):
        super().__init__(name=name)
        self._params = tree.map_structure_with_path(create_variable, params)
        self._apply = jax2tf.convert(lambda p, x: apply_fn(p, None, x))
        self._apply = tf.autograph.experimental.do_not_convert(self._apply)

    def __call__(self, inputs):
        return self._apply(self._params, inputs)

net = JaxModule(params, f.apply)
[v.name for v in net.trainable_variables]
```

```
[4]: ['jax_module/mlp/_/linear_0/b:0',
      'jax_module/mlp/_/linear_0/w:0',
      'jax_module/mlp/_/linear_1/b:0',
      'jax_module/mlp/_/linear_1/w:0',
      'jax_module/mlp/_/linear_2/b:0',
      'jax_module/mlp/_/linear_2/w:0']
```

1.5.3 Train using TensorFlow

TensorFlow datasets is a great library with lots of common datasets that you might want to do research with. Here we will use it to load the MNIST handwritten digit dataset and define a simple pipeline that will randomly shuffle training images and normalize them into $[0, 1)$.

```
[5]: import tensorflow_datasets as tfds

ds_train, ds_test = tfds.load('mnist', split=('train', 'test'),
                              shuffle_files=True, as_supervised=True)

def normalize_img(image, label):
    """Normalizes images: `uint8` -> `float32`."""
    image = tf.cast(image, tf.float32) / 255.
    return image, label

ds_train = ds_train.map(normalize_img, num_parallel_calls=tf.data.experimental.AUTOTUNE)
ds_train = ds_train.cache()
ds_train = ds_train.shuffle(60000)
ds_train = ds_train.batch(100)
ds_train = ds_train.repeat()
ds_train = ds_train.prefetch(tf.data.experimental.AUTOTUNE)

ds_test = ds_test.map(normalize_img, num_parallel_calls=tf.data.experimental.AUTOTUNE)
ds_test = ds_test.batch(100)
ds_test = ds_test.cache()
ds_test = ds_test.prefetch(tf.data.experimental.AUTOTUNE)
```

In order to train our model we need a training loop that updates model parameters based on gradients for some loss. For this example we will use the Adam optimizer from Sonnet and perform a gradient update to our parameters for each mini-batch.

```
[6]: net = JaxModule(params, f.apply)
opt = snt.optimizers.Adam(1e-3)

@tf.function(experimental_compile=True, autograph=False)
def train_step(images, labels):
    """Performs one optimizer step on a single mini-batch."""
    with tf.GradientTape() as tape:
        images = snt.flatten(images)
        logits = net(images)
        loss = tf.nn.sparse_softmax_cross_entropy_with_logits(logits=logits,
                                                                labels=labels)

        loss = tf.reduce_mean(loss)
        params = tape.watched_variables()
```

(continues on next page)

(continued from previous page)

```

    loss += 1e-4 * sum(map(tf.nn.l2_loss, params))

    grads = tape.gradient(loss, params)
    opt.apply(grads, params)
    return loss

for step, (images, labels) in enumerate(ds_train.take(6001)):
    loss = train_step(images, labels)
    if step % 1000 == 0:
        print(f"Step {step}: {loss.numpy()}")

```

```

Step 0: 2.309901475906372
Step 1000: 0.23313118517398834
Step 2000: 0.058662284165620804
Step 3000: 0.060427404940128326
Step 4000: 0.07748399674892426
Step 5000: 0.07069656997919083
Step 6000: 0.03870276361703873

```

To evaluate how our newly trained model performs we can use top-1 accuracy on our test set.

```

[7]: def accuracy(model):
    total = 0
    correct = 0
    for images, labels in ds_test:
        predictions = tf.argmax(model(snt.flatten(images)), axis=1)
        correct += tf.math.count_nonzero(tf.equal(predictions, labels))
        total += images.shape[0]

    print("Got %d/%d (%.02f%%) correct" % (correct, total, correct / total * 100.))

```

```
accuracy(net)
```

```
Got 9805/10000 (98.05%) correct
```

It is useful to visualize predictions the model is making against the input we are providing. This can be particularly useful where the model mispredicts the label, you can see that in some cases the handwriting is a bit dubious!

```

[8]: import matplotlib.pyplot as plt

def sample(correct, rows, cols):
    """Utility function to show a sample of images."""
    n = 0

    f, ax = plt.subplots(rows, cols)
    if rows > 1:
        ax = tf.nest.flatten([tuple(ax[i]) for i in range(rows)])
    f.set_figwidth(14)
    f.set_figheight(4 * rows)

    for images, labels in ds_test:
        predictions = tf.argmax(net(snt.flatten(images)), axis=1)
        eq = tf.equal(predictions, labels)

```

(continues on next page)

(continued from previous page)

```

for i, x in enumerate(eq):
    if x.numpy() == correct:
        label = labels[i]
        prediction = predictions[i]
        image = tf.squeeze(images[i])


        ax[n].imshow(image)
        ax[n].set_title("Prediction: {}\nActual: {}".format(prediction, label))

        n += 1
        if n == (rows * cols):
            break

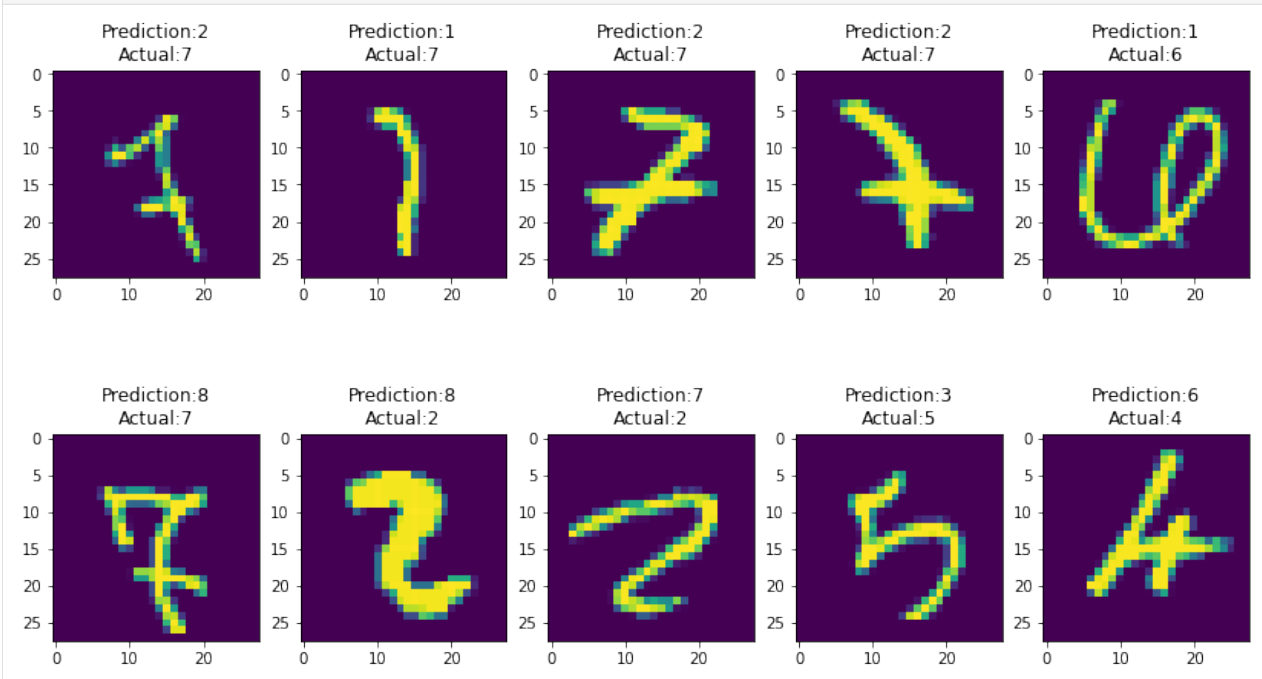
if n == (rows * cols):
    break

```

[9]: `sample(correct=True, rows=1, cols=5)`



[10]: `sample(correct=False, rows=2, cols=5)`



1.5.4 Save to disk as a TensorFlow SavedModel

It is very common to take a model trained using TensorFlow and save it to disk as a “saved model”. This is a language independent format that allows you to load your model code using Python, C++ or other languages supported by TensorFlow.

Saving

In order to save our model to disk, we need to define what the functions are that we want to save, and provide references to any state we want to save:

```
[11]: @tf.function(autograph=False, input_signature=[tf.TensorSpec([100, 28 * 28])])
def forward(x):
    return net(x)
```

```
to_save = tf.Module()
to_save.forward = forward
to_save.params = list(net.variables)
tf.saved_model.save(to_save, "/tmp/example_saved_model")
```

```
INFO:tensorflow:Assets written to: /tmp/example_saved_model/assets
```

```
INFO:tensorflow:Assets written to: /tmp/example_saved_model/assets
```

Loading

Loading a saved model is trivial, and you can see that this looks a lot like the model we saved:

```
[12]: loaded = tf.saved_model.load("/tmp/example_saved_model")
preds = loaded.forward(tf.ones([100, 28 * 28]))
assert preds.shape == [100, 10]
assert len(loaded.params) == 6
```

```
[v.name for v in loaded.params]
```

```
WARNING:tensorflow:Importing a function (__inference_forward_26770) with ops with custom_
↳gradients. Will likely fail if a gradient is requested.
```

```
WARNING:tensorflow:Importing a function (__inference_forward_26770) with ops with custom_
↳gradients. Will likely fail if a gradient is requested.
```

```
[12]: ['jax_module/mlp/_/linear_0/b:0',
'jax_module/mlp/_/linear_0/w:0',
'jax_module/mlp/_/linear_1/b:0',
'jax_module/mlp/_/linear_1/w:0',
'jax_module/mlp/_/linear_2/b:0',
'jax_module/mlp/_/linear_2/w:0']
```

Thankfully the restored model performs just as well as the model that we saved:

```
[13]: accuracy(loaded.forward)
```

```
Got 9805/10000 (98.05%) correct
```

1.6 Build your own Haiku

In this Colab, we will build a highly-simplified version of Haiku from scratch, to give you some insight into how Haiku works.

This is an “advanced” tutorial for folks seeking a deeper understanding of Haiku’s internals. It’s not required to understand how to use Haiku in practice. (“advanced” is in quotes because it’s not actually all that complicated, so don’t be afraid!)

The implementation here is based on the design of the real Haiku library, but with most details simplified. Therefore, while this should give you a reasonably accurate sense of what’s going on under-the-hood conceptually, don’t rely on the details to match.

1.6.1 The Problem

We want to be able to write object-oriented classes with parameter attributes, like this,

```
[1]: class MyModule:
      def apply(self, x):
          return self.w * x
```

and automatically transform them into pure functions, like this:

```
[2]: def my_stateless_apply(params, x):
      return params['w'] * x
```

(However, instead of using attribute access via `self.*`, we define our own accessor function called `get_param()`. It makes it much easier to intercept its usage, which we need to collect and inject parameter values later.)

Additionally, it would be nice if this transformation also defined parameter initialisation, and automatically handled assigning parameters unique names, as managing that manually in large networks can get unwieldy. E.g., if some other module also called its parameter `w`, we’d like to automatically resolve such conflicts.

We will tackle this problem in steps.

- At the first step, we will implement a basic `transform` that converts object-oriented-style functions into pure ones.
- The next step will be to add the initialisation.
- Finally, we will handle the plumbing involved when several copies of the same module are used, or different modules use the same name for their parameters.

At that stage, we will already be able to define and train a simple neural network just like with the real Haiku.

1.6.2 The Basic Strategy

We will define a function that implements the transformation from the stateful style that uses `get_param` to a stateless function. This function will be aptly called `transform`. It will wrap a `MyModule().apply` into a function that works just like `my_stateless_apply`.

Here’s how it will work. `transform(f)` will return a wrapped version of `f` that accepts an extra `params` argument. When called, it will run `f`, and every time `f` will call `get_param`, it will extract the corresponding value from `params` and return it.

```
[3]: # Global state which holds the parameters for the transformed function.
# get_param uses this to know where to get params from.
current_params = []

def transform(f):

    def apply_f(params, *args, **kwargs):
        current_params.append(params)
        outs = f(*args, **kwargs)
        current_params.pop()
        return outs

    return apply_f

def get_param(identifier):
    return current_params[-1][identifier]
```

let's test it:

```
[4]: params = dict(w=5)
my_stateless_apply(params, 5)
```

```
[4]: 25
```

```
[5]: class MyModule:
    def apply(self, x):
        return get_param('w') * x

transform(MyModule().apply)(params, 5)
```

```
[5]: 25
```

“Hold on!” you say. Isn’t JAX all about not having global state? This won’t possibly work in JAX! Well, let’s try it with JAX:

```
[6]: import jax
import jax.numpy as jnp

def linear(x):
    return x @ get_param('w') + get_param('b')

params = dict(w=jnp.ones((3, 5)), b=jnp.ones((5,)))
apply = transform(linear)

jax.jit(apply)(params, jnp.ones((10, 3)))
```

```
WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and
↪rerun for more info.)
```

```
[6]: DeviceArray([[4., 4., 4., 4., 4.],
                 [4., 4., 4., 4., 4.],
                 [4., 4., 4., 4., 4.],
                 [4., 4., 4., 4., 4.],
                 [4., 4., 4., 4., 4.]])
```

(continues on next page)

(continued from previous page)

```
[4., 4., 4., 4., 4.],
[4., 4., 4., 4., 4.],
[4., 4., 4., 4., 4.],
[4., 4., 4., 4., 4.],
[4., 4., 4., 4., 4.]], dtype=float32)
```

The reason this works is that while we use global state, we're careful about how we use it. We leave the global state after a function call the same as before it, and we ensure that the outputs of a wrapped function depend only on its inputs. Thus, JAX is none the wiser – for all it cares, the transformed function is pure.

1.6.3 Adding initialization

So far, so good, but we aren't able to reuse modules, because our transform will share parameters between all copies of the same module, because they will all be named the same. Also, defining the initial state is a pain – can we automate that?

Let's tackle initialisation first. For simplicity, our parameters will always initialise using the normal distribution, but it's not hard to add the option of different initialisers.

In this new version, we transform one object-oriented stateful function into two pure functions: one which initialises params, and one which applies them. These correspond to running the original function in two modes: initialisation and application.

To support this, we add extra machinery (the `Frame`) to track which mode we're in, and change the behaviour of `get_param()`:

- We add a `shape` argument, which tells us what shape the param should be if initialising.
- If initialising, `get_param()` will create the param of the correct shape and add it to the current params in the `Frame` before returning.

Thus, `get_param()` generates the initial values just in time for them to get used in a call of the stateful function.

```
[7]: from typing import NamedTuple, Dict, Callable
import numpy as np
```

```
[8]: # Since we're tracking more than just the current params,
# we introduce the concept of a frame as the object that holds
# state during a transformed execution.
frame_stack = []

class Frame(NamedTuple):
    """Tracks what's going on during a call of a transformed function."""
    params: Dict[str, jax.Array]
    is_initialising: bool = False

def current_frame():
    return frame_stack[-1]

class Transformed(NamedTuple):
    init: Callable
    apply: Callable
```

(continues on next page)

(continued from previous page)

```

def transform(f) -> Transformed:

    def init_f(*args, **kwargs):
        frame_stack.append(Frame({}, is_initialising=True))
        f(*args, **kwargs)
        frame = frame_stack.pop()
        return frame.params

    def apply_f(params, *args, **kwargs):
        frame_stack.append(Frame(params))
        outs = f(*args, **kwargs)
        frame_stack.pop()
        return outs

    return Transformed(init_f, apply_f)

def get_param(identifier, shape):
    if current_frame().is_initialising:
        current_frame().params[identifier] = np.random.normal(size=shape)

    return current_frame().params[identifier]

```

Let's test it by implementing a Linear module:

```

[9]: # Make printing parameters a little more readable
def parameter_shapes(params):
    return jax.tree_util.tree_map(lambda p: p.shape, params)

class Linear:

    def __init__(self, width):
        self._width = width

    def __call__(self, x):
        w = get_param('w', shape=(x.shape[-1], self._width))
        b = get_param('b', shape=(self._width,))
        return x @ w + b

init, apply = transform(Linear(4))

data = jnp.ones((2, 3))

params = init(data)
parameter_shapes(params)

[9]: {'b': (4,), 'w': (3, 4)}

```

```
[10]: apply(params, data)
```

```
[10]: DeviceArray([[ -1.0345883,  0.3280404, -2.4382973,  0.5717376],
                  [ -1.0345883,  0.3280404, -2.4382973,  0.5717376]],
                dtype=float32)
```

1.6.4 Adding unique parameter names: our finished mini-Haiku

Alright! Time to tackle nesting modules, and our prototype will be done.

For this, we need to give each parameter an unambiguous name. Here, we will use a scheme that's somewhat different and incompatible with the real Haiku, but which is simpler. The idea is to record the names of the functions being called, and to assign each parameter a unique identifier based on its location in the call stack.

For this, we will define a Module class. Each module will have a unique identifier based on the class name and the number of instances of the module created so far. (Real Haiku allows to customise these names, but we ignore that for simplicity)

We will also define a decorator for Module methods, called `module_method`, which will tell us when the wrapped function is called, allowing us to track the current parameter scope. Real haiku uses metaclasses to automatically wrap all methods on a Module, but for simplicity we do this manually.

```
[11]: import dataclasses
import collections

@dataclasses.dataclass
class Frame:
    """Tracks what's going on during a call of a transformed function."""
    params: Dict[str, jax.Array]
    is_initialising: bool = False

    # Keeps track of how many modules of each class have been created so far.
    # Used to assign new modules unique names.
    module_counts: Dict[str, int] = dataclasses.field(
        default_factory=lambda: collections.defaultdict(lambda: 0))

    # Keeps track of the entire path to the current module method call.
    # Module methods, when called, will add themselves to this stack.
    # Used to give each parameter a unique name corresponding to the
    # method scope it is in.
    call_stack: list = dataclasses.field(default_factory=list)

    def create_param_path(self, identifier) -> str:
        """Creates a unique path for this param."""
        return '/'.join(['~'] + self.call_stack + [identifier])

    def create_unique_module_name(self, module_name: str) -> str:
        """Assigns a unique name to the module by appending its number to its name."""
        number = self.module_counts[module_name]
        self.module_counts[module_name] += 1
        return f"{module_name}_{number}"

frame_stack = []

def current_frame():
    return frame_stack[-1]

class Module:
    def __init__(self):
        # Assign a unique (for the current `transform` call)
```

(continues on next page)

(continued from previous page)

```

# name to this instance of the module.
self._unique_name = current_frame().create_unique_module_name(
    self.__class__.__name__)

def module_method(f):
    """A decorator for Module methods."""
    # In the real Haiku, this doesn't face the user but is applied by a metaclass.

    def wrapped(self, *args, **kwargs):
        """A version of f that lets the frame know it's being called."""
        # Self is the instance to which this method is attached.
        module_name = self._unique_name
        call_stack = current_frame().call_stack
        call_stack.append(module_name)
        call_stack.append(f.__name__)
        outs = f(self, *args, **kwargs)
        assert call_stack.pop() == f.__name__
        assert call_stack.pop() == module_name
        return outs

    return wrapped

def get_param(identifier, shape):
    frame = current_frame()
    param_path = frame.create_param_path(identifier)

    if frame.is_initialising:
        frame.params[param_path] = np.random.normal(size=shape)

    return frame.params[param_path]

class Linear(Module):

    def __init__(self, width):
        super().__init__()
        self._width = width

    @module_method # Again, this decorator is behind-the-scenes in real Haiku.
    def __call__(self, x):
        w = get_param('w', shape=(x.shape[-1], self._width))
        b = get_param('b', shape=(self._width,))
        return x @ w + b

```

At this stage, we have replicated some core Haiku functionality, but we still don't have: * control over initialisation * rng handling * state handling (though conceptually that's analogous to parameter handling) * any kind of validation and error handling * freezing parameters once they're created * thread-safety * JAX transformations inside a transform (e.g. `hk.remat`) * JAX control flow inside transforms (e.g. `hk.scan`) * last but not least, documentation :)

and lots more. However, the basics work, so we can take our mini-Haiku for a ride:

```
[12]: init, apply = transform(lambda x: Linear(4)(x))
```

```
params = init(data)
parameter_shapes(params)
```

```
[12]: {'~/Linear_0/___call___/b': (4,), '~/Linear_0/___call___/w': (3, 4)}
```

```
[13]: apply(params, data)
```

```
[13]: DeviceArray([[ -1.1969297,  1.3215988,  5.175427 , -1.9018829],
                 [ -1.1969297,  1.3215988,  5.175427 , -1.9018829]]),
          dtype=float32)
```

Different Modules in a function call all have separate parameters:

```
[14]: class MLP(Module):
```

```
    def __init__(self, widths):
        super().__init__()
        self._widths = widths
```

```
    @module_method
```

```
    def __call__(self, x):
        for w in self._widths:
            out = Linear(w)(x)
            x = jax.nn.sigmoid(out)
        return out
```

```
[15]: init, apply = transform(lambda x: MLP([3, 5])(x))
parameter_shapes(init(data))
```

```
[15]: {'~/MLP_0/___call___/Linear_0/___call___/b': (3,),
      '~/MLP_0/___call___/Linear_0/___call___/w': (3, 3),
      '~/MLP_0/___call___/Linear_1/___call___/b': (5,),
      '~/MLP_0/___call___/Linear_1/___call___/w': (3, 5)}
```

While the same module called in different places reuses parameters:

```
[16]: class ParameterReuseTest(Module):
```

```
    @module_method
```

```
    def __call__(self, x):
        f = Linear(x.shape[-1])

        x = f(x)
        x = jax.nn.relu(x)
        return f(x)
```

```
init, forward = transform(lambda x: ParameterReuseTest()(x))
parameter_shapes(init(data))
```

```
[16]: {'~/ParameterReuseTest_0/___call___/Linear_0/___call___/b': (3,),
      '~/ParameterReuseTest_0/___call___/Linear_0/___call___/w': (3, 3)}
```

1.6.5 Example training loop

```
[17]: import matplotlib.pyplot as plt

[18]: # Data: a quadratic curve.
xs = np.linspace(-2., 2., num=128)[: , None] # Generate array of shape (128, 1).
ys = xs ** 2

# Model
def mlp(x):
    return MLP([128, 128, 1])(x)

init, forward = transform(mlp)
params = init(xs)
parameter_shapes(params)

[18]: {'~/MLP_0/___call___/Linear_0/___call___/b': (128,),
 '~/MLP_0/___call___/Linear_0/___call___/w': (1, 128),
 '~/MLP_0/___call___/Linear_1/___call___/b': (128,),
 '~/MLP_0/___call___/Linear_1/___call___/w': (128, 128),
 '~/MLP_0/___call___/Linear_2/___call___/b': (1,),
 '~/MLP_0/___call___/Linear_2/___call___/w': (128, 1)}


[19]: # Loss function and update function
def loss_fn(params, x, y):
    return jnp.mean((forward(params, x) - y) ** 2)

LEARNING_RATE = 0.003

@jax.jit
def update(params, x, y):
    grads = jax.grad(loss_fn)(params, x, y)
    return jax.tree_util.tree_map(
        lambda p, g: p - LEARNING_RATE * g, params, grads
    )

[20]: for _ in range(5000):
    params = update(params, xs, ys)

[21]: plt.scatter(xs, ys, label='Data')
plt.scatter(xs, forward(params, xs), label='Model prediction')
plt.legend()
plt.show()
```



[]:

1.7 Visualization

Haiku supports two ways to visualize your program. To use these you need to install two additional dependencies:

```
[1]: !pip install dm-tree graphviz
```

```
Requirement already satisfied: dm-tree in /tmp/haiku-env/lib/python3.11/site-packages (0.
↪1.8)
Requirement already satisfied: graphviz in /tmp/haiku-env/lib/python3.11/site-packages.
↪(0.20.1)
```

```
[2]: import jax
import jax.numpy as jnp
import haiku as hk
```

1.7.1 Tabulate

Like many neural network libraries, Haiku supports showing a summary of the execution of your program as a table of modules. Haiku's approach is to trace the execution of your program and to produce a table of (interesting) module method calls.

For example, the interesting methods for a 3 layer MLP would be `MLP.__call__` which in turns calls `Linear.__call__` on three inner modules. For each module method we show columns relating to the input/output size of arrays, as well as details of the modules parameters and where it fits in the module hierarchy.

```
[3]: def f(x):
    return hk.nets.MLP([300, 100, 10])(x)

f = hk.transform(f)
x = jnp.ones([8, 28 * 28])

print(hk.experimental.tabulate(f)(x))
```

No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and rerun for more info.)

Module	Config	Module params
Input	Param count	Param bytes
mlp (MLP)	MLP(output_sizes=[300, 100, 10])	
f32[8,784] f32[8,10]	266,610 1.07 MB	
mlp/~linear_0 (Linear)	Linear(output_size=300, name='linear_0')	w: f32[784,300]
f32[8,784] f32[8,300]	235,500 942.00 KB	
mlp (MLP)		b: f32[300]
mlp/~linear_1 (Linear)	Linear(output_size=100, name='linear_1')	w: f32[300,100]
f32[8,300] f32[8,100]	30,100 120.40 KB	
mlp (MLP)		b: f32[100]
mlp/~linear_2 (Linear)	Linear(output_size=10, name='linear_2')	w: f32[100,10]
f32[8,100] f32[8,10]	1,010 4.04 KB	
mlp (MLP)		b: f32[10]

We also offer access to the raw data used to build this table if you want to create your own summary:

```
[4]: for method_invocation in hk.experimental.eval_summary(f)(x):
    print(method_invocation)
```

```
MethodInvocation(module_details=ModuleDetails(module=MLP(output_sizes=[300, 100, 10]),
↳ method_name='__call__', params={'mlp/~linear_0/b': f32[300], 'mlp/~linear_0/w':
↳ f32[784,300], 'mlp/~linear_1/b': f32[100], 'mlp/~linear_1/w': f32[300,100], 'mlp/~
↳ linear_2/b': f32[10], 'mlp/~linear_2/w': f32[100,10]}, state={}), args_spec=(f32[8,
↳ 784],), kwargs_spec={}, output_spec=f32[8,10], context=MethodContext(module=MLP(output
↳ sizes=[300, 100, 10]), method_name='__call__', orig_method=functools.partial(<function
↳ MLP.__call__ at 0x7f173d83f600>, MLP(output_sizes=[300, 100, 10])), orig_class=<class
↳ 'haiku._src.nets.mlp.MLP'>), call_stack=(ModuleDetails(module=MLP(output_sizes=[300,
↳ 100, 10]), method_name='__call__', params={'mlp/~linear_0/b': f32[300], 'mlp/~linear_
↳ 0/w': f32[784,300], 'mlp/~linear_1/b': f32[100], 'mlp/~linear_1/w': f32[300,100],
↳ 'mlp/~linear_2/b': f32[10], 'mlp/~linear_2/w': f32[100,10]}, state={})),)
MethodInvocation(module_details=ModuleDetails(module=Linear(output_size=300, name=
↳ 'linear_0'), method_name='__call__', params={'mlp/~linear_0/b': f32[300], 'mlp/~
↳ linear_0/w': f32[784,300]}, state={}), args_spec=(f32[8,784],), kwargs_spec={}, output_
↳ spec=f32[8,300], context=MethodContext(module=Linear(output_size=300, name='linear_0'),
↳ method_name='__call__', orig_method=functools.partial(<function Linear.__call__ at
↳ 0x7f173d927e20>, Linear(output_size=300, name='linear_0')), orig_class=<class 'haiku._
↳ src.basic.Linear'>), call_stack=(ModuleDetails(module=Linear(output_size=300, name=
↳ 'linear_0'), method_name='__call__', params={'mlp/~linear_0/b': f32[300], 'mlp/~
↳ linear_0/w': f32[784,300]}, state={}), ModuleDetails(module=MLP(output_sizes=[300, 100,
↳ 10]), method_name='__call__', params={'mlp/~linear_0/b': f32[300], 'mlp/~linear_0/w
↳ f32[784,300], 'mlp/~linear_1/b': f32[100], 'mlp/~linear_1/w': f32[300,100], 'mlp/~
↳ linear_2/b': f32[10], 'mlp/~linear_2/w': f32[100,10]}, state={})))
```

(continued from previous page)

```

MethodInvocation(module_details=ModuleDetails(module=Linear(output_size=100, name=
↳ 'linear_1'), method_name='__call__', params={'mlp/~/'linear_1/b': f32[100], 'mlp/~/'
↳ linear_1/w': f32[300,100]}, state={}), args_spec=(f32[8,300]), kwargs_spec={}, output_
↳ spec=f32[8,100], context=MethodContext(module=Linear(output_size=100, name='linear_1'),
↳ method_name='__call__', orig_method=functools.partial(<function Linear.__call__ at
↳ 0x7f173d927e20>, Linear(output_size=100, name='linear_1')), orig_class=<class 'haiku._
↳ src.basic.Linear'>), call_stack=(ModuleDetails(module=Linear(output_size=100, name=
↳ 'linear_1'), method_name='__call__', params={'mlp/~/'linear_1/b': f32[100], 'mlp/~/'
↳ linear_1/w': f32[300,100]}, state={}), ModuleDetails(module=MLP(output_sizes=[300, 100,
↳ 10]), method_name='__call__', params={'mlp/~/'linear_0/b': f32[300], 'mlp/~/'linear_0/w
↳ ': f32[784,300], 'mlp/~/'linear_1/b': f32[100], 'mlp/~/'linear_1/w': f32[300,100], 'mlp/~
↳ /linear_2/b': f32[10], 'mlp/~/'linear_2/w': f32[100,10]}, state={})))
MethodInvocation(module_details=ModuleDetails(module=Linear(output_size=10, name='linear_
↳ 2'), method_name='__call__', params={'mlp/~/'linear_2/b': f32[10], 'mlp/~/'linear_2/w':
↳ f32[100,10]}, state={}), args_spec=(f32[8,100]), kwargs_spec={}, output_spec=f32[8,
↳ 10], context=MethodContext(module=Linear(output_size=10, name='linear_2'), method_name=
↳ '__call__', orig_method=functools.partial(<function Linear.__call__ at 0x7f173d927e20>,
↳ Linear(output_size=10, name='linear_2')), orig_class=<class 'haiku._src.basic.Linear'>
↳ ), call_stack=(ModuleDetails(module=Linear(output_size=10, name='linear_2'), method_
↳ name='__call__', params={'mlp/~/'linear_2/b': f32[10], 'mlp/~/'linear_2/w': f32[100,10]},
↳ state={}), ModuleDetails(module=MLP(output_sizes=[300, 100, 10]), method_name='__call_
↳ _', params={'mlp/~/'linear_0/b': f32[300], 'mlp/~/'linear_0/w': f32[784,300], 'mlp/~/'
↳ linear_1/b': f32[100], 'mlp/~/'linear_1/w': f32[300,100], 'mlp/~/'linear_2/b': f32[10],
↳ 'mlp/~/'linear_2/w': f32[100,10]}, state={})))

```

1.7.2 Graphviz (aka. to_dot)

Haiku supports rendering your program as a graphviz graph. We show all of the JAX primitives involved in a given computation clustered by Haiku module.

Lets start by visualizing a simple program not using Haiku modules:

```

[5]: def f(a):
      b = jnp.sin(a)
      c = jnp.cos(b)
      d = b + c
      e = a + d
      return e

x = jnp.ones([1])
dot = hk.to_dot(f)(x)

import graphviz
graphviz.Source(dot)

```

[5]: The visualization above shows our program as a simple dataflow graph of our single input highlighted in orange (args[0]) being passed through some operations and producing a result (highlighted in blue). Primitive operations (e.g. sin, cos and add) are highlighted in yellow.

Actual Haiku programs are often far more complex, involving many modules and many more primitive operations. For these programs it is often useful to visualize the program on a module by module basis.

to_dot offers this by clustering operations by their module. Again it is probably simplest to see an example:

```
[6]: def f(x):
      return hk.nets.MLP([300, 100, 10])(x)

f = hk.transform(f)

rng = jax.random.PRNGKey(42)
x = jnp.ones([8, 28 * 28])
params = f.init(rng, x)

dot = hk.to_dot(f.apply)(params, None, x)
graphviz.Source(dot)
```

```
[6]:
```

1.8 Training a subset of parameters

Sometimes when training a neural network it is useful to hold some parameters of your network fixed while updating others. This is commonly referred to as “non-trainable variables” or “layer freezing”.

In typical neural network training, parameters are updated by computing gradients and computing an update via an optimizer such as SGD or ADAM. Updates are then applied to parameters and the process repeats until you have converged.

As such to implement “layer freezing” or “non-trainable variables” in JAX, we simply need to not compute and apply updates for certain parameters of our network.

In JAX computing gradients and applying updates to parameters are fully in your control as a user. JAX’s autodiff mechanics allow you to compute gradients wrt any positional argument to a function.

In Haiku (and other NN libraries) it is typical to pass your parameters as a single positional argument to your function (e.g. `grads = jax.grad(loss_fn)(params, ...)`).

To support taking gradients wrt a subset of parameters, we need to allow users to split their parameters into two positional arguments, such that they can compute gradients wrt a subset of their parameters (e.g. `trainable_params_grads = jax.grad(loss_fn)(trainable_params, non_trainable_params, ...)`).

Haiku ships with some utilities that make it easier to manipulate the parameters dictionary in order to split into these trainable/non-trainable sets as well as to recombine your parameters into a single dictionary.

We will walk through how to do this with a simple MLP and teach it the identity function.

```
[ ]: import haiku as hk
      import jax
      import jax.numpy as jnp
      import numpy as np
```

The forward pass of our network is a standard MLP. We want to adjust the parameters of this MLP such that it computes the identity. That is `forward([[1.0], [2.0], [3.0]]) == [1, 2, 3]`. We will do this for a maximum of 10 numbers.

Our network starts randomly initialised so the results initially do not make much sense:

```
[ ]: num_classes = 10

def f(x):
    return hk.nets.MLP([300, 100, num_classes])(x)
```

(continues on next page)

(continued from previous page)

```
f = hk.transform(f)

def test(params, num_classes=num_classes):
    x = np.arange(num_classes).reshape([num_classes, 1]).astype(np.float32)
    y = jnp.argmax(f.apply(params, None, x), axis=-1)
    for x, y in zip(x, y):
        print(x, "->", y)

rng = jax.random.PRNGKey(42)
x = np.zeros([num_classes, 1])
params = f.init(rng, x)

print("before training")
test(params)
```

```
WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and
↳rerun for more info.)
/tmp/haiku-docs-env/lib/python3.8/site-packages/jax/_src/lax/lax.py:6271: UserWarning:
↳Explicitly requested dtype float64 requested in zeros is not available, and will be
↳truncated to dtype float32. To enable more dtypes, set the jax_enable_x64
↳configuration option or the JAX_ENABLE_X64 shell environment variable. See https://
↳github.com/google/jax#current-gotchas for more.
warnings.warn(msg.format(dtype, fun_name , truncated_dtype))
```

```
before training
[0.] -> 0
[1.] -> 3
[2.] -> 3
[3.] -> 3
[4.] -> 3
[5.] -> 3
[6.] -> 3
[7.] -> 3
[8.] -> 3
[9.] -> 3
```

It is useful to visualise our parameters so we can compare to their final state:

```
[ ]: import matplotlib.pyplot as plt
import seaborn as sns

sns.set_theme()

def plot_params(params):
    fig, axs = plt.subplots(ncols=2, nrows=3)
    fig.tight_layout()
    fig.set_figwidth(12)
    fig.set_figheight(6)
    for row, module in enumerate(sorted(params)):
        ax = axs[row][0]
        sns.heatmap(params[module]["w"], cmap="YlGnBu", ax=ax)
        ax.title.set_text(f"{module}/w")
```

(continues on next page)

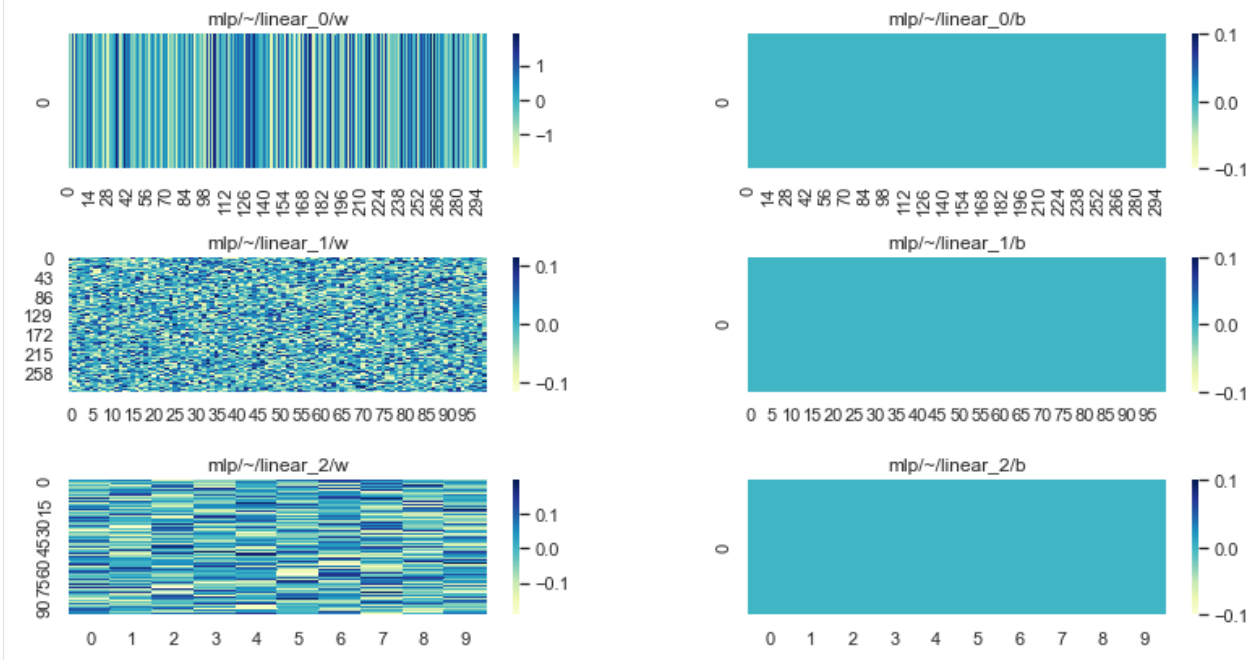
(continued from previous page)

```

ax = axs[row][1]
b = np.expand_dims(params[module][“b”], axis=0)
sns.heatmap(b, cmap=“YlGnBu”, ax=ax)
ax.title.set_text(f“{module}/b”)

```

```
plot_params(params)
```



To train our network we’ll create some simple synthetic batches of data:

```

[ ]: def dataset(*, batch_size, num_records):
    for _ in range(num_records):
        y = np.arange(num_classes)
        y = np.random.permutation(y)[:batch_size]
        x = y.reshape([batch_size, 1]).astype(np.float32)
        yield x, y

for x, y in dataset(batch_size=4, num_records=5):
    print("x :=", x.tolist(), "y :=", y)

```

```

x := [[0.0], [8.0], [7.0], [1.0]] y := [0 8 7 1]
x := [[6.0], [7.0], [0.0], [9.0]] y := [6 7 0 9]
x := [[4.0], [0.0], [9.0], [6.0]] y := [4 0 9 6]
x := [[0.0], [4.0], [6.0], [5.0]] y := [0 4 6 5]
x := [[4.0], [3.0], [0.0], [5.0]] y := [4 3 0 5]

```

Now for the interesting part. Lets pretend that we only want to update the parameters of the first and last layer of our MLP.

The simplest and most efficient way to do this is to partition our parameters into two groups, “trainable” and “non trainable”. Haiku provides a convenience function for doing this in `hk.data_structures.partition(...)`:

```
[ ]: # Partition our params into trainable and non trainable explicitly.
trainable_params, non_trainable_params = hk.data_structures.partition(
    lambda m, n, p: m != "mlp/~linear_1", params)

print("trainable:", list(trainable_params))
print("non_trainable:", list(non_trainable_params))

trainable: ['mlp/~linear_0', 'mlp/~linear_2']
non_trainable: ['mlp/~linear_1']
```

The reason we split our parameters is that this allows us to pass them to our loss function as separate positional arguments.

In JAX gradients are taken with respect to positional arguments. By splitting our parameters into two groups we can take gradients with respect to just one of the positional arguments. We can then use those gradients to update a subset of our parameters.

The last piece of the puzzle is that we need to combine our “trainable” and “non trainable” parameters together before calling our apply function. Again Haiku provides `hk.data_structures.merge(...)` to make this easy:

```
[ ]: def loss_fn(trainable_params, non_trainable_params, images, labels):
    # NOTE: We need to combine trainable and non trainable before calling apply.
    params = hk.data_structures.merge(trainable_params, non_trainable_params)

    # NOTE: From here on this is a standard softmax cross entropy loss.
    logits = f.apply(params, None, images)
    labels = jax.nn.one_hot(labels, logits.shape[-1])
    return -jnp.sum(labels * jax.nn.log_softmax(logits)) / labels.shape[0]

def sgd_step(params, grads, *, lr):
    return jax.tree_util.tree_map(lambda p, g: p - g * lr, params, grads)

def train_step(trainable_params, non_trainable_params, x, y):
    # NOTE: We will only compute gradients wrt `trainable_params`.
    trainable_params_grads = jax.grad(loss_fn)(trainable_params,
                                                non_trainable_params, x, y)

    # NOTE: We are only updating `trainable_params`.
    trainable_params = sgd_step(trainable_params, trainable_params_grads, lr=0.1)
    return trainable_params

train_step = jax.jit(train_step)

for x, y in dataset(batch_size=num_classes, num_records=10000):
    # NOTE: In our training loop only our trainable parameters are updated.
    trainable_params = train_step(trainable_params, non_trainable_params, x, y)
```

We can see that even though we only trained a subset of our parameters, our NN is able to learn this simple function:

```
[ ]: # Merge params again for inference.
params = hk.data_structures.merge(trainable_params, non_trainable_params)

print("after training")
test(params)
```

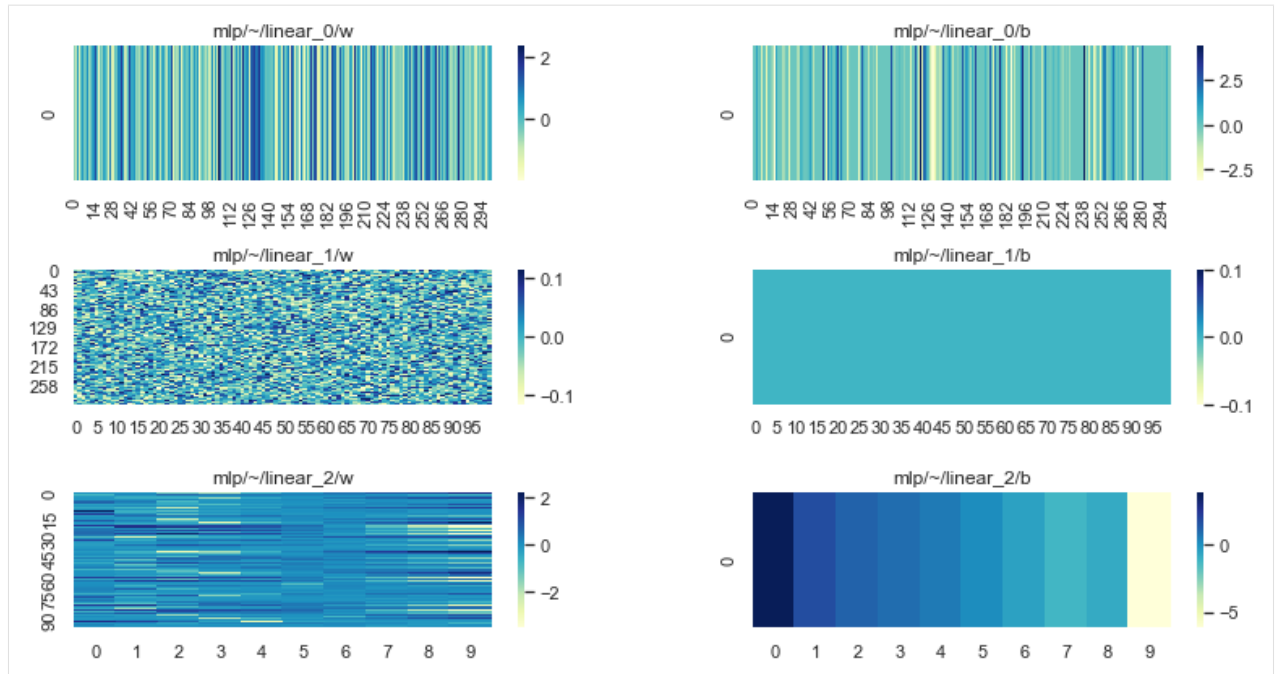
```
after training
[0.] -> 0
[1.] -> 1
[2.] -> 2
[3.] -> 3
[4.] -> 4
[5.] -> 5
[6.] -> 6
[7.] -> 7
[8.] -> 8
[9.] -> 9
```

Of course it is not smart enough to generalize to out of distribution inputs:

```
[ ]: test(params, num_classes=num_classes+10)
[0.] -> 0
[1.] -> 1
[2.] -> 2
[3.] -> 3
[4.] -> 4
[5.] -> 5
[6.] -> 6
[7.] -> 7
[8.] -> 8
[9.] -> 9
[10.] -> 9
[11.] -> 9
[12.] -> 9
[13.] -> 9
[14.] -> 9
[15.] -> 9
[16.] -> 9
[17.] -> 9
[18.] -> 9
[19.] -> 9
```

Looking at our parameters we can see that `linear_1` is still in its initial state (randomly initialised weight matrix and zero initialized bias):

```
[ ]: plot_params(params)
```



1.8.1 Freezing layers with Optax

Alternatively, Optax users can fix parameters using ``optax.multi_transform`` <https://optax.readthedocs.io/en/latest/api.html#optax.multi_transform>`__`. Users can read more [here](#).

1.9 Parameter sharing in Haiku

1.9.1 Introduction

In Haiku, parameter reuse is determined uniquely by *module instance names*, i.e., if a module instance has the same name as another module instance, they share parameters.

Unless specified, module names are automatically determined by Haiku based on the module *class* name (following a pattern that was established in TensorFlow 1 with Sonnet V1). More in detail, module naming follows these rules:

1. Module names are assigned when the module instance is *constructed*. Unless a module instance name is provided as an argument to the constructor, Haiku generates one from the current *module class name* (basically: `to_snake_case(CurrentClassName)`).
2. If the module instance name doesn't end in a `_N` (where `N` is a number) and another module instance with the same name already exists, Haiku adds an incremental number to the end of the new module instance name (e.g. `module_1`).
3. When two modules are nested (i.e., a module instance is constructed inside another module's class definition), then the inner module name will be prepended by the *outer module name* and, possibly (see the next point), the *outer module current method* being called. The constructor (i.e., `__init__`) is replaced by the tilde `~` symbol.
4. If the calling method name is `__call__` this will be ignored (the method name will be prepended by the *outer module name* only).

- When there are multiple layers of nesting, the previous rule is applied at each level of nesting, and each inner module name is based on the module name and calling method name of the module immediately preceding the current module in the hierarchy of calls.

Let's see how this works with a practical example.

1.9.2 Flat modules (no nesting)

This section covers parameter sharing when the modules are not nested.

```
[4]: ##@title Imports and accessory functions
import functools
import haiku as hk
import jax
import jax.numpy as jnp

def parameter_shapes(params):
    """Make printing parameters a little more readable."""
    return jax.tree_util.tree_map(lambda p: p.shape, params)

def transform_and_print_shapes(fn, x_shape=(2, 3)):
    """Print name and shape of the parameters."""
    rng = jax.random.PRNGKey(42)
    x = jnp.ones(x_shape)

    transformed_fn = hk.transform(fn)
    params = transformed_fn.init(rng, x)
    print('\nThe name and shape of the parameters are:')
    print(parameter_shapes(params))

def assert_all_equal(params_1, params_2):
    assert all(jax.tree_util.tree_leaves(
        jax.tree_util.tree_map(lambda a, b: (a == b).all(), params_1, params_2)))
```

```
[6]: w_init = hk.initializers.TruncatedNormal(stddev=1)

class SimpleModule(hk.Module):
    """A simple module class with one variable."""

    def __init__(self, output_channels, name=None):
        super().__init__(name)
        assert isinstance(output_channels, int)
        self._output_channels = output_channels

    def __call__(self, x):
        w_shape = (x.shape[-1], self._output_channels)
        w = hk.get_parameter("w", w_shape, x.dtype, init=w_init)
        return jnp.dot(x, w)
```

```
[ ]: def f(x):
    # This instance will be named `a_simple_module`.
```

(continues on next page)

(continued from previous page)

```

simple = SimpleModule(output_channels=2)
simple_out = simple(x) # implicitly calls module_install.__call__()
print(f'The name assigned to "simple" is: "{simple.module_name}".')
return simple_out

transform_and_print_shapes(f)

```

WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and rerun for more info.)

The name assigned to "simple" is: "simple_module".

The name and shape of the parameters are:
{'simple_module': {'w': (3, 2)}}

Great! Here we see that indeed if we create a `SimpleModule` instance and do not specify a name, Haiku assigns to it the name `a_simple_module`. This is also reflected in the parameters associated to the module.

What happens if we instantiate `SimpleModule` twice though? Does Haiku assign to both instances the same name?

```

[ ]: def f(x):
    # This instance will be named `a_simple_module`.
    simple_one = SimpleModule(output_channels=2)
    # This instance will be named `a_simple_module_1`.
    simple_two = SimpleModule(output_channels=2)
    first_out = simple_one(x)
    second_out = simple_two(x)
    print(f'The name assigned to "simple_one" is: "{simple_one.module_name}".')
    print(f'The name assigned to "simple_two" is: "{simple_two.module_name}".')
    return first_out + second_out

transform_and_print_shapes(f)

```

The name assigned to "simple_one" is: "simple_module".
The name assigned to "simple_two" is: "simple_module_1".

The name and shape of the parameters are:
{'simple_module': {'w': (3, 2)}, 'simple_module_1': {'w': (3, 2)}}

As expected Haiku is smart enough to differentiate the two instances and avoid accidental parameter sharing: the second instance is named `a_simple_module_1` and each instance has its own set of parameters. Good!

But what if we wanted to share parameters? In this case, we would have to instantiate the module only once and *call* it multiple times. Let's see how this works:

```

[ ]: def f(x):
    # This instance will be named `a_simple_module`.
    simple_one = SimpleModule(output_channels=2)
    first_out = simple_one(x)
    second_out = simple_one(x) # share parameters w/ previous call
    print(f'The name assigned to "simple_one" is: "{simple_one.module_name}".')
    return first_out + second_out

transform_and_print_shapes(f)

```

The name assigned to "simple_one" is: "simple_module".

The name and shape of the parameters are:

```
{'simple_module': {'w': (3, 2)}}
```

1.9.3 Nested modules

In this section we'll see what happens when we nest one `hk.Module` into another.

```
[ ]: class NestedModule(hk.Module):
    """A module class with a nested module created in the constructor."""

    def __init__(self, output_channels, name=None):
        super().__init__(name)
        assert isinstance(output_channels, int)
        self._output_channels = output_channels
        self.inner_simple = SimpleModule(self._output_channels)

    def __call__(self, x):
        w_shape = (x.shape[-1], self._output_channels)
        # Another variable that is also called `w`.
        w = hk.get_parameter("w", w_shape, x.dtype, init=w_init)
        return jnp.dot(x, w) + self.inner_simple(x)

[ ]: def f(x):
    # This will be named `a_nested_module` and the SimpleModule instance created
    # inside it will be named `a_nested_module/a_simple_module`.
    nested = NestedModule(output_channels=2)
    nested_out = nested(x)
    print('The name assigned to outer module (i.e., "nested") is: '
          f'{nested.module_name}'.)
    print('The name assigned to the inner module (i.e., inside "nested") is: '
          f'{nested.inner_simple.module_name}'.)
    return nested_out
```

```
transform_and_print_shapes(f)
```

The name assigned to outer module (i.e., "nested") is: "nested_module".

The name assigned to the inner module (i.e., inside "nested") is: "nested_module/~/
↪simple_module".

The name and shape of the parameters are:

```
{'nested_module': {'w': (3, 2)}, 'nested_module/~/'
  'simple_module': {'w': (3, 2)}}
```

As expected, the inner module name depends on: (a) the outer module name; and (b) the outer module's method being called.

Note also how the outer module's constructor name `__init__` is replaced by a `~` in the parameter names. If the inner module instance was created inside the `__call__` method of the outer module, the inner module instance name would have been `'a_nested_module/a_simple_module'`.

In this example we defined all the modules from scratch, but the same holds for any of the modules and networks defined in Haiku, e.g., `hk.Linear`, `hk.nets.MLP`, If you are curious, see what happens if you assign to `self.inner_simple` an instance of `hk.Linear` instead of `SimpleModule`.

Let's try now multiple levels of nesting:

```
[ ]: class TwiceNestedModule(hk.Module):
    """A module class with a nested module containing a nested module."""

    def __init__(self, output_channels, name=None):
        super().__init__(name)
        assert isinstance(output_channels, int)
        self._output_channels = output_channels
        self.inner_nested = NestedModule(self._output_channels)

    def __call__(self, x):
        w_shape = (x.shape[-1], self._output_channels)
        w = hk.get_parameter("w", w_shape, x.dtype, init=w_init)
        return jnp.dot(x, w) + self.inner_nested(x)

[ ]: def f(x):
    """Create the module instances and inspect their names."""
    # Instantiate a NestedModule instance. This will be named `a_nested_module`.
    # The SimpleModule instance created inside it will be named
    # a_nested_module/a_simple_module`.
    outer = TwiceNestedModule(output_channels=2)
    outer_out = outer(x)
    print(f'The name assigned to the most outer class is: "{outer.module_name}".')
    print('The name assigned to the module inside "double_nested" is: "'
          f'{outer.inner_nested.module_name}'.')
    print('The name assigned to the module inside it is "'
          f'{outer.inner_nested.inner_simple.module_name}'.')
    return outer_out
```

```
transform_and_print_shapes(f)
```

```
The name assigned to the most outer class is: "twice_nested_module".
The name assigned to the module inside "double_nested" is: "twice_nested_module/~/_nested_
↪module".
The name assigned to the module inside it is "twice_nested_module/~/_nested_module/~/_
↪simple_module".
```

The name and shape of the parameters are:

```
{'twice_nested_module': {'w': (3, 2)}, 'twice_nested_module/~/_nested_module': {'w': (3,
↪2)}, 'twice_nested_module/~/_nested_module/~/_simple_module': {'w': (3, 2)}}
```

Great, this also works as expected: the full hierarchy of module names and calls is reflected in the inner module names.

1.9.4 Multitransform: merge the parameters without sharing them

Sometimes when we have multiple transformed functions it can be convenient to merge all the parameters in a unique structure, to reduce the number of dictionaries we have to store and pass around. It can be the case though that some of these functions instantiate the same modules, and we want to make sure that their parameters don't get shared accidentally.

`hk.multi_transform` comes to rescue in this case, and merges the parameters in a unique dictionary making sure that duplicated parameters are renamed to avoid accidental sharing.

```
[ ]: def f(x):
    """A SimpleModule followed by a Linear layer."""
    module_instance = SimpleModule(output_channels=2)
    out = module_instance(x)
    linear = hk.Linear(40)
    return linear(out)

def g(x):
    """A SimpleModule followed by an MLP."""
    module_instance = SimpleModule(output_channels=2)
    return module_instance(x) * 2 # twice

# Transform both functions, and print their respective parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))
transformed_f = hk.transform(f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(g)
params_g = transformed_g.init(rng, x)
print('f parameters:', parameter_shapes(params_f))
print('g parameters:', parameter_shapes(params_g))

# Transform both functions at once with hk.multi_transform, and print the
# resulting merged parameter structure.

def multitransform_f_and_g():
    def template(x):
        return f(x), g(x)
    return template, (f, g)
init, (f_apply, g_apply) = hk.multi_transform(multitransform_f_and_g)
merged_params = init(rng, x)

print('\nThe name and shape of the multi-transform parameters are:\n',
      parameter_shapes(merged_params))

f parameters: {'linear': {'b': (40,), 'w': (2, 40)}, 'simple_module': {'w': (3, 2)}}
g parameters: {'simple_module': {'w': (3, 2)}}

The name and shape of the multi-transform parameters are:
{'linear': {'b': (40,), 'w': (2, 40)}, 'simple_module': {'w': (3, 2)}, 'simple_module_1
↪': {'w': (3, 2)}}
```

In this example `f` and `g` both instantiate a `SimpleModule` instance with the same arguments, and if we transform them separately we see that both dictionaries contain a `'simple_module'` key.

When we transform them together instead, `hk.multi_transform` takes care of us of renaming one of them to

'simple_module_1', thus preventing accidental parameter sharing.

1.9.5 Sharing parameters between transformed functions

Now that we understood how module names are assigned and how this affects parameter sharing, let's see how we can share parameters between transformed functions.

In this section we will consider two functions, `f` and `g`, and explore different strategies to share parameters. We will consider a number of cases that differ in how many of the modules instantiated by each function are the same, and if their parameters have the same shape.

Case 1: All modules have the same names, and the same shape

Let's reuse one of the modules we created before, and try to instantiate it twice inside two different functions:

```
[ ]: def f(x):
    """Apply SimpleModule to x."""
    module_instance = SimpleModule(output_channels=2)
    out = module_instance(x)
    return out

def g(x):
    """Like f, but double the output"""
    module_instance = SimpleModule(output_channels=2)
    out = module_instance(x)
    return out * 2

# Transform both functions, and print the parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))

transformed_f = hk.transform(f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(g)
params_g = transformed_g.init(rng, x)

print('f parameters:', parameter_shapes(params_f))
print('g parameters:', parameter_shapes(params_g))

f parameters: {'simple_module': {'w': (3, 2)}}
g parameters: {'simple_module': {'w': (3, 2)}}
```

Great! Since `f` and `g` are using exactly the same modules, the sets of initialized variables generated with each have the same *name structure* (note that the actual values might differ, depending on initialization).

Now, if we wanted to share parameters in this case, we could initialize only one of the two functions (e.g., `f`) and use the resulting parameters for both functions, i.e., when we call `transformed_f.apply` and `transformed_g.apply`.

Case 2: Common modules have the same names, and the same shape

This is a nice trick, but what if the functions were not identical? Let's build two such functions:

```
[ ]: def f(x):
    """A SimpleModule followed by a Linear layer."""
    module_instance = SimpleModule(output_channels=2)
    out = module_instance(x)
    linear = hk.Linear(40)
    return linear(out)

def g(x):
    """A SimpleModule followed by an MLP."""
    module_instance = SimpleModule(output_channels=2)
    out = module_instance(x)
    linear = hk.nets.MLP((10, 40))
    return linear(out)

# Transform both functions, and print the parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))

transformed_f = hk.transform(f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(g)
params_g = transformed_g.init(rng, x)

print('\nThe name and shape of the f parameters are:\n',
      parameter_shapes(params_f))
print('\nThe name and shape of the g parameters are:\n',
      parameter_shapes(params_g))
```

```
The name and shape of the f parameters are:
{'linear': {'b': (40,), 'w': (2, 40)}, 'simple_module': {'w': (3, 2)}}
```

```
The name and shape of the g parameters are:
{'mlp~/~/linear_0': {'b': (10,), 'w': (2, 10)}, 'mlp~/~/linear_1': {'b': (40,), 'w': (10, 40)}, 'simple_module': {'w': (3, 2)}}
```

Now we have a problem! Both sets of parameters have a 'simple_module' component, but they also each contain parameters that are specific only to that function, so we cannot simply initialise only one of the functions and use the returned parameters for both as we did before. But we would still like to share the parameters of 'simple_module'. How can we do that?

One option here is to use `haiku.data_structures.merge` <https://dm-haiku.readthedocs.io/en/latest/api.html#haiku.data_structures.merge>`__` to combine the two sets of parameters. This will merge the two structures, keeping only the value from the last structure when both structures have the same parameters (i.e., 'simple_module' in our example). Let try that:

```
[ ]: merged_params = hk.data_structures.merge(params_f, params_g)
print('\nThe name and shape of the shared parameters are:\n',
      parameter_shapes(merged_params))
```

The name and shape of the shared parameters are:

```
{'linear': {'b': (40,), 'w': (2, 40)}, 'mlp/~linear_0': {'b': (10,), 'w': (2, 10)},
↪ 'mlp/~linear_1': {'b': (40,), 'w': (10, 40)}, 'simple_module': {'w': (3, 2)}}
```

Brilliant! Now we have a shared set of parameters that contains all the disjoint parameters and a single set of parameters for the shared 'simple_module'. Let's verify that we can use this set of parameters when calling either function:

```
[ ]: f_out = transformed_f.apply(merged_params, rng, x)
      g_out = transformed_g.apply(merged_params, rng, x)
```

```
print('f_out mean:', f_out.mean())
print('g_out mean:', g_out.mean())
```

```
f_out mean: 0.037986994
g_out mean: 0.104857825
```

This gives us little control over what gets shared though: what if the two functions had parameters with the same name that we don't want to share?

Case 3: Common modules have the same names, but different shapes

Let's modify our previous example to use a `hk.Linear` layer in both functions:

```
[ ]: def f(x):
      """A SimpleModule followed by two Linear layers."""
      module_instance = SimpleModule(output_channels=2)
      out = module_instance(x)
      mlp = hk.nets.MLP((10, 5))
      out = mlp(out)
      last_linear = hk.Linear(4)
      return last_linear(out)

      def g(x):
          """Same as f, with a bigger final layer."""
          module_instance = SimpleModule(output_channels=2)
          out = module_instance(x)
          mlp = hk.nets.MLP((10, 5))
          out = mlp(out)
          last_linear = hk.Linear(20) # another Linear, but bigger
          return last_linear(out)

      # Transform both functions, and print the parameter shapes.
      rng = jax.random.PRNGKey(42)
      x = jnp.ones((2, 3))

      transformed_f = hk.transform(f)
      params_f = transformed_f.init(rng, x)
      transformed_g = hk.transform(g)
      params_g = transformed_g.init(rng, x)

      print('\nThe name and shape of the f parameters are:\n',
            parameter_shapes(params_f))
```

(continues on next page)

(continued from previous page)

```
print('\n\nThe name and shape of the g parameters are:\n',
      parameter_shapes(params_g))
```

The name and shape of the f parameters are:

```
{'linear': {'b': (4,), 'w': (5, 4)}, 'mlp/~linear_0': {'b': (10,), 'w': (2, 10)}, 'mlp/
↳~/linear_1': {'b': (5,), 'w': (10, 5)}, 'simple_module': {'w': (3, 2)}}
```

The name and shape of the g parameters are:

```
{'linear': {'b': (20,), 'w': (5, 20)}, 'mlp/~linear_0': {'b': (10,), 'w': (2, 10)},
↳'mlp/~linear_1': {'b': (5,), 'w': (10, 5)}, 'simple_module': {'w': (3, 2)}}
```

Now we have a problem! Both sets of parameters have a 'linear' component, but their respective parameters have different shapes. If we merged them as we did before, the parameters of the 'linear' from f would be dropped and we couldn't use the merged parameters to call it:

```
merged_params = hk.data_structures.merge(params_f, params_g)
print('\n\nThe name and shape of the merged parameters are:\n',
      parameter_shapes(merged_params))
```

```
f_out = transformed_f.apply(merged_params, rng, x) # fails
# ValueError: 'linear/w' with retrieved shape (5, 20) does not match shape=[5, 4]
↳dtype=dtype('float32')
```

How can we share the parameters of 'simple_module' and mlp, but keep the parameters of the two output linear layers separated?

A solution would be to instantiate `simple_module` and `mlp` outside of the functions, so that they get instantiated only once, and then use that instance in both functions. But all Haiku modules must be initialised in a transform, so doing so naively would incur in an error:

```
module_instance = SimpleModule(output_channels=2) # this fails
# ValueError: All `hk.Module`s must be initialized inside an `hk.transform`.
mlp = hk.nets.MLP((10, 5))

def f(x):
    """A SimpleModule followed by a Linear layer."""
    out = module_instance(x)
    out = mlp(out)
    linear = hk.Linear(4)
    return linear(out)

def g(x):
    """A SimpleModule followed by a bigger Linear layer."""
    out = module_instance(x)
    out = mlp(out)
    linear = hk.Linear(20) # another Linear, but bigger
    return linear(out)
```

We can work around that by creating another function

```
[ ]: class CachedModule():
    def __call__(self, *inputs):
```

(continues on next page)

(continued from previous page)

```

# Create the instances if are not in the cache.
if not hasattr(self, 'cached_simple_module'):
    self.cached_simple_module = SimpleModule(output_channels=2)
if not hasattr(self, 'cached_mlp'):
    self.cached_mlp = hk.nets.MLP((10, 5))

# Apply the cached instances.
out = self.cached_simple_module(*inputs)
out = self.cached_mlp(out)
return out

def f(x):
    """A SimpleModule followed by a Linear layer."""
    shared_preprocessing = CachedModule()
    out = shared_preprocessing(x)
    linear = hk.Linear(4)
    return linear(out)

def g(x):
    """A SimpleModule followed by a bigger Linear layer."""
    shared_preprocessing = CachedModule()
    out = shared_preprocessing(x)
    linear = hk.Linear(20) # another Linear, but bigger
    return linear(out)

# Transform both functions, and print the parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))

transformed_f = hk.transform(f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(g)
params_g = transformed_g.init(rng, x)

print('\nThe name and shape of the f parameters are:\n',
      parameter_shapes(params_f))
print('\nThe name and shape of the g parameters are:\n',
      parameter_shapes(params_g))

# Verify that the simple module parameters are shared.
assert_all_equal(params_f['mlp~/linear_0'],
                 params_g['mlp~/linear_0'])
assert_all_equal(params_f['mlp~/linear_1'],
                 params_g['mlp~/linear_1'])
print('\nThe MLP parameters are shared!')

```

```

The name and shape of the f parameters are:
{'linear': {'b': (4,), 'w': (5, 4)}, 'mlp~/linear_0': {'b': (10,), 'w': (2, 10)}, 'mlp/
~~/linear_1': {'b': (5,), 'w': (10, 5)}, 'simple_module': {'w': (3, 2)}}

```

(continues on next page)

(continued from previous page)

The name and shape of the g parameters are:

```
{'linear': {'b': (20,), 'w': (5, 20)}, 'mlp/~linear_0': {'b': (10,), 'w': (2, 10)},
↪ 'mlp/~linear_1': {'b': (5,), 'w': (10, 5)}, 'simple_module': {'w': (3, 2)}}
```

The MLP parameters are shared!

If we want to share a big number of modules it can become tedious to cache each one of them manually inside of `CachedModule`. Furthermore, it would be nice if we didn't have to define a different `CachedModule` object for every function we want to cache.

We can use `hk.to_module` to create a more general `CachedModule` object that takes an arbitrary Haiku function and caches it:

```
[ ]: class CachedModule():
    """Cache one instance of the function and call it multiple times."""
    def __init__(self, fn):
        self._fn = fn

    def __call__(self, *args, **kwargs):
        if not hasattr(self, "_instance"):
            ModularisedFn = hk.to_module(self._fn)
            self._instance = ModularisedFn()
        return self._instance(*args, **kwargs)

def shared_preprocessing_fn(x):
    simple_module = SimpleModule(output_channels=2)
    out = simple_module(x)
    mlp = hk.nets.MLP((10, 5))
    return mlp(out)

def f(x):
    """A SimpleModule followed by a Linear layer."""
    shared_preprocessing = CachedModule(shared_preprocessing_fn)
    out = shared_preprocessing(x)
    linear = hk.Linear(4)
    return linear(out)

def g(x):
    """A SimpleModule followed by a bigger Linear layer."""
    shared_preprocessing = CachedModule(shared_preprocessing_fn)
    out = shared_preprocessing(x)
    linear = hk.Linear(20) # another Linear, but bigger
    return linear(out)

# Transform both functions, and print the parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))

transformed_f = hk.transform(f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(g)
```

(continues on next page)

(continued from previous page)

```

params_g = transformed_g.init(rng, x)

print('\nThe name and shape of the f parameters are:\n',
      parameter_shapes(params_f))
print('\nThe name and shape of the g parameters are:\n',
      parameter_shapes(params_g))

# Verify that the simple module parameters are shared.
assert_all_equal(params_f['shared_preprocessing_fn/mlp/~linear_0'],
                 params_g['shared_preprocessing_fn/mlp/~linear_0'])
assert_all_equal(params_f['shared_preprocessing_fn/mlp/~linear_1'],
                 params_g['shared_preprocessing_fn/mlp/~linear_1'])
print('\nThe MLP parameters are shared!')

```

The name and shape of the f parameters are:

```

{'linear': {'b': (4,), 'w': (5, 4)}, 'shared_preprocessing_fn/mlp/~linear_0': {'b': (10,), 'w': (2, 10)}, 'shared_preprocessing_fn/mlp/~linear_1': {'b': (5,), 'w': (10, 5)}, 'shared_preprocessing_fn/simple_module': {'w': (3, 2)}}

```

The name and shape of the g parameters are:

```

{'linear': {'b': (20,), 'w': (5, 20)}, 'shared_preprocessing_fn/mlp/~linear_0': {'b': (10,), 'w': (2, 10)}, 'shared_preprocessing_fn/mlp/~linear_1': {'b': (5,), 'w': (10, 5)}, 'shared_preprocessing_fn/simple_module': {'w': (3, 2)}}

```

The MLP parameters are shared!

When we work with objects it can also be convenient to define a decorator to do the same:

```

[7]: def share_parameters():
      def decorator(fn):
          def wrapper(*args, **kwargs):
              if wrapper.instance is None:
                  wrapper.instance = hk.to_module(fn)()
              return wrapper.instance(*args, **kwargs)
          wrapper.instance = None
          return functools.wraps(fn)(wrapper)
      return decorator

class Wrapper():

    @share_parameters()
    def shared_preprocessing(self, x):
        simple_module = SimpleModule(output_channels=2)
        out = simple_module(x)
        mlp = hk.nets.MLP((10, 5))
        return mlp(out)

    def f(self, x):
        """A SimpleModule followed by a Linear layer."""
        out = self.shared_preprocessing(x)
        linear = hk.Linear(4)

```

(continues on next page)

(continued from previous page)

```

    return linear(out)

def g(self, x):
    """A SimpleModule followed by a bigger Linear layer."""
    out = self.shared_preprocessing(x)
    linear = hk.Linear(20) # another Linear, but bigger
    return linear(out)

# Transform both functions, and print the parameter shapes.
rng = jax.random.PRNGKey(42)
x = jnp.ones((2, 3))

wrapper = Wrapper()
transformed_f = hk.transform(wrapper.f)
params_f = transformed_f.init(rng, x)
transformed_g = hk.transform(wrapper.g)
params_g = transformed_g.init(rng, x)

print('\nThe name and shape of the f parameters are:\n',
      parameter_shapes(params_f))
print('\nThe name and shape of the g parameters are:\n',
      parameter_shapes(params_g))

# Verify that the simple module parameters are shared.
assert_all_equal(params_f['shared_preprocessing/mlp/~linear_0'],
                 params_g['shared_preprocessing/mlp/~linear_0'])
assert_all_equal(params_f['shared_preprocessing/mlp/~linear_1'],
                 params_g['shared_preprocessing/mlp/~linear_1'])
print('\nThe MLP parameters are shared!')

```

The name and shape of the f parameters are:

```

{'linear': {'b': (4,), 'w': (5, 4)}, 'shared_preprocessing/mlp/~linear_0': {'b': (10,),
↪ 'w': (2, 10)}, 'shared_preprocessing/mlp/~linear_1': {'b': (5,), 'w': (10, 5)},
↪ 'shared_preprocessing/simple_module': {'w': (3, 2)}}

```

The name and shape of the g parameters are:

```

{'linear': {'b': (20,), 'w': (5, 20)}, 'shared_preprocessing/mlp/~linear_0': {'b': (10,
↪ ), 'w': (2, 10)}, 'shared_preprocessing/mlp/~linear_1': {'b': (5,), 'w': (10, 5)},
↪ 'shared_preprocessing/simple_module': {'w': (3, 2)}}

```

The MLP parameters are shared!

KNOWN ISSUES

Warning: Using JAX transformations like `jax.jit()` and `jax.remat()` inside of Haiku networks can lead to hard to interpret tracing errors and potentially silently wrong results. Read *Limitations of Nesting JAX Functions and Haiku Modules* to find out how to work around these issues.

CONTRIBUTE

- Issue tracker: <https://github.com/deepmind/dm-haiku/issues>
- Source code: <https://github.com/deepmind/dm-haiku/tree/main>

SUPPORT

If you are having issues, please let us know by filing an issue on our [issue tracker](#).

LICENSE

Haiku is licensed under the Apache 2.0 License.

INDICES AND TABLES

- genindex
- modindex

BIBLIOGRAPHY

- [1] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014. URL: <https://arxiv.org/abs/1409.2329>.
- [2] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. An empirical exploration of recurrent network architectures. In *International Conference on Machine Learning*, 2342–2350. 2015.
- [3] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: a machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, 802–810. 2015.

PYTHON MODULE INDEX

h

- haiku, 133
- haiku.config, 124
- haiku.data_structures, 125
- haiku.experimental, 113
- haiku.experimental.flax, 32
- haiku.experimental.jaxpr_info, 121
- haiku.initializers, 89
- haiku.mixed_precision, 111
- haiku.nets, 96
- haiku.pad, 93
- haiku.testing, 130

Symbols

- `__call__` () (*haiku.AvgPool* method), 51
- `__call__` () (*haiku.BatchApply* method), 86
- `__call__` () (*haiku.BatchNorm* method), 66
- `__call__` () (*haiku.Bias* method), 50
- `__call__` () (*haiku.ConvND* method), 54
- `__call__` () (*haiku.ConvNDTranspose* method), 58
- `__call__` () (*haiku.Deferred* method), 88
- `__call__` () (*haiku.EMAParamsTree* method), 72
- `__call__` () (*haiku.Embed* method), 87
- `__call__` () (*haiku.ExponentialMovingAverage* method), 71
- `__call__` () (*haiku.GRU* method), 78
- `__call__` () (*haiku.GroupNorm* method), 67
- `__call__` () (*haiku.IdentityCore* method), 80
- `__call__` () (*haiku.LSTM* method), 77
- `__call__` () (*haiku.LayerNorm* method), 69
- `__call__` () (*haiku.Linear* method), 49
- `__call__` () (*haiku.MaxPool* method), 52
- `__call__` () (*haiku.MultiHeadAttention* method), 83
- `__call__` () (*haiku.RMSNorm* method), 70
- `__call__` () (*haiku.RNNCore* method), 73
- `__call__` () (*haiku.ResetCore* method), 79
- `__call__` () (*haiku.Reshape* method), 85
- `__call__` () (*haiku.SNParamsTree* method), 72
- `__call__` () (*haiku.SeparableDepthwiseConv2D* method), 64
- `__call__` () (*haiku.Sequential* method), 53
- `__call__` () (*haiku.SpectralNorm* method), 70
- `__call__` () (*haiku.VanillaRNN* method), 76
- `__call__` () (*haiku.initializers.Constant* method), 89
- `__call__` () (*haiku.initializers.Identity* method), 90
- `__call__` () (*haiku.initializers.Orthogonal* method), 90
- `__call__` () (*haiku.initializers.RandomNormal* method), 91
- `__call__` () (*haiku.initializers.RandomUniform* method), 91
- `__call__` () (*haiku.initializers.TruncatedNormal* method), 91
- `__call__` () (*haiku.initializers.UniformScaling* method), 93
- `__call__` () (*haiku.initializers.VarianceScaling* method), 92
- `__call__` () (*haiku.nets.MLP* method), 96
- `__call__` () (*haiku.nets.MobileNetV1* method), 97
- `__call__` () (*haiku.nets.ResNet* method), 99
- `__call__` () (*haiku.nets.ResNet.BlockGroup* method), 98
- `__call__` () (*haiku.nets.ResNet.BlockV1* method), 98
- `__call__` () (*haiku.nets.ResNet.BlockV2* method), 98
- `__call__` () (*haiku.nets.VectorQuantizer* method), 104
- `__call__` () (*haiku.nets.VectorQuantizerEMA* method), 105
- `__delattr__` () (*haiku.Deferred* method), 89
- `__delattr__` () (*haiku.experimental.ArraySpec* method), 116
- `__delattr__` () (*haiku.experimental.MethodInvocation* method), 117
- `__delattr__` () (*haiku.experimental.ModuleDetails* method), 118
- `__eq__` () (*haiku.experimental.ArraySpec* method), 116
- `__eq__` () (*haiku.experimental.MethodInvocation* method), 117
- `__eq__` () (*haiku.experimental.ModuleDetails* method), 118
- `__hash__` () (*haiku.experimental.ArraySpec* method), 116
- `__hash__` () (*haiku.experimental.MethodInvocation* method), 117
- `__hash__` () (*haiku.experimental.ModuleDetails* method), 118
- `__init__` () (*haiku.AvgPool* method), 51
- `__init__` () (*haiku.BatchApply* method), 86
- `__init__` () (*haiku.BatchNorm* method), 65
- `__init__` () (*haiku.Bias* method), 49
- `__init__` () (*haiku.Conv1D* method), 55
- `__init__` () (*haiku.Conv1DLSTM* method), 81
- `__init__` () (*haiku.Conv1DTranspose* method), 59
- `__init__` () (*haiku.Conv2D* method), 56
- `__init__` () (*haiku.Conv2DLSTM* method), 81
- `__init__` () (*haiku.Conv2DTranspose* method), 59
- `__init__` () (*haiku.Conv3D* method), 57
- `__init__` () (*haiku.Conv3DLSTM* method), 82
- `__init__` () (*haiku.Conv3DTranspose* method), 60
- `__init__` () (*haiku.ConvND* method), 54

- `__init__` () (*haiku.ConvNDTranspose method*), 58
 - `__init__` () (*haiku.DeepRNN method*), 78
 - `__init__` () (*haiku.Deferred method*), 88
 - `__init__` () (*haiku.DepthwiseConv1D method*), 61
 - `__init__` () (*haiku.DepthwiseConv2D method*), 62
 - `__init__` () (*haiku.DepthwiseConv3D method*), 62
 - `__init__` () (*haiku.EMAParamsTree method*), 72
 - `__init__` () (*haiku.Embed method*), 86
 - `__init__` () (*haiku.ExponentialMovingAverage method*), 71
 - `__init__` () (*haiku.Flatten method*), 85
 - `__init__` () (*haiku.GRU method*), 77
 - `__init__` () (*haiku.GroupNorm method*), 66
 - `__init__` () (*haiku.InstanceNorm method*), 67
 - `__init__` () (*haiku.LSTM method*), 77
 - `__init__` () (*haiku.LayerNorm method*), 68
 - `__init__` () (*haiku.Linear method*), 48
 - `__init__` () (*haiku.MaxPool method*), 51
 - `__init__` () (*haiku.Module method*), 17
 - `__init__` () (*haiku.MultiHeadAttention method*), 82
 - `__init__` () (*haiku.PRNGSequence method*), 26
 - `__init__` () (*haiku.RMSNORM method*), 69
 - `__init__` () (*haiku.ResetCore method*), 79
 - `__init__` () (*haiku.Reshape method*), 84
 - `__init__` () (*haiku.SNParamsTree method*), 71
 - `__init__` () (*haiku.SeparableDepthwiseConv2D method*), 63
 - `__init__` () (*haiku.Sequential method*), 53
 - `__init__` () (*haiku.SpectralNorm method*), 70
 - `__init__` () (*haiku.VanillaRNN method*), 76
 - `__init__` () (*haiku.experimental.ArraySpec method*), 116
 - `__init__` () (*haiku.experimental.MethodInvocation method*), 117
 - `__init__` () (*haiku.experimental.ModuleDetails method*), 118
 - `__init__` () (*haiku.initializers.Constant method*), 89
 - `__init__` () (*haiku.initializers.Identity method*), 90
 - `__init__` () (*haiku.initializers.Orthogonal method*), 90
 - `__init__` () (*haiku.initializers.RandomNormal method*), 90
 - `__init__` () (*haiku.initializers.RandomUniform method*), 91
 - `__init__` () (*haiku.initializers.TruncatedNormal method*), 91
 - `__init__` () (*haiku.initializers.UniformScaling method*), 93
 - `__init__` () (*haiku.initializers.VarianceScaling method*), 92
 - `__init__` () (*haiku.nets.MLP method*), 96
 - `__init__` () (*haiku.nets.MobileNetV1 method*), 97
 - `__init__` () (*haiku.nets.ResNet method*), 99
 - `__init__` () (*haiku.nets.ResNet.BlockGroup method*), 98
 - `__init__` () (*haiku.nets.ResNet.BlockV1 method*), 98
 - `__init__` () (*haiku.nets.ResNet.BlockV2 method*), 99
 - `__init__` () (*haiku.nets.ResNet101 method*), 101
 - `__init__` () (*haiku.nets.ResNet152 method*), 102
 - `__init__` () (*haiku.nets.ResNet18 method*), 100
 - `__init__` () (*haiku.nets.ResNet200 method*), 102
 - `__init__` () (*haiku.nets.ResNet34 method*), 100
 - `__init__` () (*haiku.nets.ResNet50 method*), 101
 - `__init__` () (*haiku.nets.VectorQuantizer method*), 103
 - `__init__` () (*haiku.nets.VectorQuantizerEMA method*), 105
 - `__next__` () (*haiku.PRNGSequence method*), 27
 - `__post_init__` () (*haiku.Module method*), 18
 - `__setattr__` () (*haiku.Deferred method*), 89
 - `__setattr__` () (*haiku.experimental.ArraySpec method*), 116
 - `__setattr__` () (*haiku.experimental.MethodInvocation method*), 117
 - `__setattr__` () (*haiku.experimental.ModuleDetails method*), 118
- ## A
- `abstract_to_dot` () (*in module haiku.experimental*), 113
 - `apply` (*haiku.MultiTransformed attribute*), 31
 - `apply` (*haiku.MultiTransformedWithState attribute*), 31
 - `apply` (*haiku.Transformed attribute*), 30
 - `apply` (*haiku.TransformedWithState attribute*), 30
 - `args_spec` (*haiku.experimental.MethodInvocation attribute*), 117
 - `ARRAY_INDEX` (*haiku.EmbedLookupStyle attribute*), 87
 - `ArraySpec` (*class in haiku.experimental*), 116
 - `as_html` () (*in module haiku.experimental.jaxpr_info*), 122
 - `as_html_page` () (*in module haiku.experimental.jaxpr_info*), 123
 - `avg_pool` () (*in module haiku*), 50
 - `AvgPool` (*class in haiku*), 50
- ## B
- `BatchApply` (*class in haiku*), 86
 - `BatchNorm` (*class in haiku*), 65
 - `Bias` (*class in haiku*), 49
- ## C
- `call_stack` (*haiku.experimental.MethodInvocation attribute*), 117
 - `causal` () (*in module haiku.pad*), 95
 - `cell` (*haiku.LSTMState attribute*), 29
 - `check_jax_usage` () (*in module haiku.experimental*), 118
 - `clear_policy` () (*in module haiku.mixed_precision*), 113
 - `commitment_cost` (*haiku.nets.VectorQuantizer attribute*), 103

- commitment_cost (*haiku.nets.VectorQuantizerEMA* attribute), 105
 cond() (*in module haiku*), 106
 Constant (*class in haiku.initializers*), 89
 context (*haiku.experimental.MethodInvocation* attribute), 117
 context() (*in module haiku.config*), 124
 Conv1D (*class in haiku*), 55
 Conv1DLSTM (*class in haiku*), 80
 Conv1DTranspose (*class in haiku*), 59
 Conv2D (*class in haiku*), 56
 Conv2DLSTM (*class in haiku*), 81
 Conv2DTranspose (*class in haiku*), 59
 Conv3D (*class in haiku*), 57
 Conv3DLSTM (*class in haiku*), 82
 Conv3DTranspose (*class in haiku*), 60
 ConvND (*class in haiku*), 54
 ConvNDTranspose (*class in haiku*), 58
 create() (*in module haiku.pad*), 94
 create_from_padfn() (*in module haiku.pad*), 94
 create_from_tuple() (*in module haiku.pad*), 95
 css() (*in module haiku.experimental.jaxpr_info*), 123
 current_name() (*in module haiku*), 41
 current_policy() (*in module haiku.mixed_precision*), 112
 custom_creator() (*in module haiku*), 21
 custom_getter() (*in module haiku*), 21
 custom_setter() (*in module haiku*), 22
- ## D
- decay (*haiku.nets.VectorQuantizerEMA* attribute), 105
 deep_rnn_with_skip_connections() (*in module haiku*), 78
 DeepRNN (*class in haiku*), 78
 Deferred (*class in haiku*), 88
 DepthwiseConv1D (*class in haiku*), 61
 DepthwiseConv2D (*class in haiku*), 62
 DepthwiseConv3D (*class in haiku*), 62
 DO_NOT_STORE (*in module haiku*), 42
 dropout() (*in module haiku*), 52
 dtype (*haiku.experimental.ArraySpec* attribute), 116
 dynamic_unroll() (*in module haiku*), 74
- ## E
- EMAParamsTree (*class in haiku*), 72
 Embed (*class in haiku*), 86
 embedding_dim (*haiku.nets.VectorQuantizer* attribute), 103
 embedding_dim (*haiku.nets.VectorQuantizerEMA* attribute), 104
 EmbedLookupStyle (*class in haiku*), 87
 epsilon (*haiku.nets.VectorQuantizerEMA* attribute), 105
 eval_shape() (*in module haiku*), 108
 eval_summary() (*in module haiku.experimental*), 116
 expand_apply() (*in module haiku*), 75
 ExponentialMovingAverage (*class in haiku*), 71
 Expression (*class in haiku.experimental.jaxpr_info*), 123
- ## F
- fast_eval_shape() (*in module haiku.experimental*), 121
 filter() (*in module haiku.data_structures*), 125
 Flatten (*class in haiku*), 85
 flatten_flax_to_haiku() (*in module haiku.experimental.flax*), 33
 force_name() (*in module haiku*), 45
 fori_loop() (*in module haiku*), 107
 format_module() (*in module haiku.experimental.jaxpr_info*), 123
 full() (*in module haiku.pad*), 95
 full_name (*haiku.GetterContext* attribute), 22
 full_name (*haiku.SetterContext* attribute), 23
- ## G
- get_channel_index() (*in module haiku*), 64
 get_current_state() (*in module haiku*), 44
 get_initial_state() (*in module haiku*), 45
 get_parameter() (*in module haiku*), 19
 get_params() (*in module haiku*), 43
 get_policy() (*in module haiku.mixed_precision*), 112
 get_state() (*in module haiku*), 19
 GetterContext (*class in haiku*), 22
 grad() (*in module haiku*), 108
 GroupNorm (*class in haiku*), 66
 GRU (*class in haiku*), 77
- ## H
- haiku
 module, 34, 132, 133
 haiku.config
 module, 124
 haiku.data_structures
 module, 125
 haiku.experimental
 module, 113
 haiku.experimental.flax
 module, 32
 haiku.experimental.jaxpr_info
 module, 121
 haiku.initializers
 module, 89
 haiku.mixed_precision
 module, 111
 haiku.nets
 module, 96
 haiku.pad

module, 93
 haiku.testing
 module, 130
 hidden (*haiku.LSTMState* attribute), 29

I

Identity (*class in haiku.initializers*), 90
 IdentityCore (*class in haiku*), 80
 init (*haiku.MultiTransformed* attribute), 31
 init (*haiku.MultiTransformedWithState* attribute), 31
 init (*haiku.Transformed* attribute), 30
 init (*haiku.TransformedWithState* attribute), 30
 initial_state() (*haiku.GRU* method), 78
 initial_state() (*haiku.IdentityCore* method), 80
 initial_state() (*haiku.LSTM* method), 77
 initial_state() (*haiku.ResetCore* method), 79
 initial_state() (*haiku.RNNCore* method), 73
 initial_state() (*haiku.VanillaRNN* method), 76
 initialize() (*haiku.ExponentialMovingAverage* method), 71
 Initializer (*in module haiku.initializers*), 89
 InstanceNorm (*class in haiku*), 67
 intercept_methods() (*in module haiku*), 24
 is_padfn() (*in module haiku.pad*), 94
 is_subset() (*in module haiku.data_structures*), 126

J

js() (*in module haiku.experimental.jaxpr_info*), 123

K

kwargs_spec (*haiku.experimental.MethodInvocation* attribute), 117

L

layer_stack (*class in haiku*), 38
 LayerNorm (*class in haiku*), 68
 LayerStackTransparencyMapping (*class in haiku*), 40
 lift() (*in module haiku*), 34
 lift() (*in module haiku.experimental.flax*), 33
 lift_with_state() (*in module haiku*), 36
 lifted_prefix_name (*haiku.GetterContext* attribute), 23
 lifted_prefix_name (*haiku.SetterContext* attribute), 23
 LiftWithStateUpdater (*class in haiku*), 38
 Linear (*class in haiku*), 48
 LSTM (*class in haiku*), 76
 LSTMState (*class in haiku*), 29

M

make_model_info() (*in module haiku.experimental.jaxpr_info*), 122
 map() (*in module haiku*), 107

map() (*in module haiku.data_structures*), 126
 max_pool() (*in module haiku*), 51
 MaxPool (*class in haiku*), 51
 maybe_get_rng_sequence_state() (*in module haiku*), 28
 maybe_next_rng_key() (*in module haiku*), 28
 merge() (*in module haiku.data_structures*), 127
 method_name (*haiku.experimental.ModuleDetails* attribute), 118
 method_name (*haiku.MethodContext* attribute), 25
 MethodContext (*class in haiku*), 25
 MethodInvocation (*class in haiku.experimental*), 117
 MLP (*class in haiku.nets*), 96
 MobileNetV1 (*class in haiku.nets*), 97
 module
 haiku, 34, 132, 133
 haiku.config, 124
 haiku.data_structures, 125
 haiku.experimental, 113
 haiku.experimental.flax, 32
 haiku.experimental.jaxpr_info, 121
 haiku.initializers, 89
 haiku.mixed_precision, 111
 haiku.nets, 96
 haiku.pad, 93
 haiku.testing, 130
 Module (*class in haiku*), 17
 Module (*class in haiku.experimental.flax*), 32
 Module (*class in haiku.experimental.jaxpr_info*), 123
 module (*haiku.experimental.ModuleDetails* attribute), 118
 module (*haiku.GetterContext* attribute), 22
 module (*haiku.MethodContext* attribute), 25
 module (*haiku.SetterContext* attribute), 23
 module_auto_repr() (*in module haiku.experimental*), 120
 module_details (*haiku.experimental.MethodInvocation* attribute), 117
 module_name (*haiku.GetterContext* attribute), 23
 module_name (*haiku.SetterContext* attribute), 24
 ModuleDetails (*class in haiku.experimental*), 118
 ModuleProtocol (*class in haiku*), 31
 multi_transform() (*in module haiku*), 13
 multi_transform_with_state() (*in module haiku*), 15
 MultiHeadAttention (*class in haiku*), 82
 multinomial() (*in module haiku*), 133
 MultiTransformed (*class in haiku*), 31
 MultiTransformedWithState (*class in haiku*), 31
 MutableParams (*in module haiku*), 30
 MutableState (*in module haiku*), 30
 N
 name (*haiku.GetterContext* attribute), 23

- name (*haiku.SetterContext* attribute), 24
 name_like() (*in module haiku*), 46
 name_scope() (*in module haiku*), 40
 next() (*haiku.PRNGSequence* method), 27
 next_rng_key() (*in module haiku*), 27
 next_rng_keys() (*in module haiku*), 27
 num_embeddings (*haiku.nets.VectorQuantizer* attribute), 103
 num_embeddings (*haiku.nets.VectorQuantizerEMA* attribute), 105
- ## O
- ONE_HOT (*haiku.EmbedLookupStyle* attribute), 87
 one_hot() (*in module haiku*), 133
 optimize_rng_use() (*in module haiku.experimental*), 119
 orig_class (*haiku.MethodContext* attribute), 26
 orig_method (*haiku.MethodContext* attribute), 26
 original_dtype (*haiku.GetterContext* attribute), 23
 original_dtype (*haiku.SetterContext* attribute), 23
 original_init (*haiku.GetterContext* attribute), 23
 original_shape (*haiku.GetterContext* attribute), 23
 original_shape (*haiku.SetterContext* attribute), 23
 Orthogonal (*class in haiku.initializers*), 90
 output_spec (*haiku.experimental.MethodInvocation* attribute), 117
- ## P
- PadFn (*in module haiku.pad*), 93
 params (*haiku.experimental.ModuleDetails* attribute), 118
 Params (*in module haiku*), 29
 params_dict() (*haiku.Module* method), 18
 partition() (*in module haiku.data_structures*), 127
 partition_n() (*in module haiku.data_structures*), 128
 PRNGSequence (*class in haiku*), 26
 push_policy() (*in module haiku.mixed_precision*), 113
- ## Q
- quantize() (*haiku.nets.VectorQuantizer* method), 104
 quantize() (*haiku.nets.VectorQuantizerEMA* method), 106
- ## R
- RandomNormal (*class in haiku.initializers*), 90
 RandomUniform (*class in haiku.initializers*), 91
 reinit() (*in module haiku*), 109
 replace_rng_sequence_state() (*in module haiku*), 29
 reserve() (*haiku.PRNGSequence* method), 26
 reserve_rng_keys() (*in module haiku*), 28
 ResetCore (*class in haiku*), 79
 Reshape (*class in haiku*), 84
 ResNet (*class in haiku.nets*), 98
 ResNet.BlockGroup (*class in haiku.nets*), 98
 ResNet.BlockV1 (*class in haiku.nets*), 98
 ResNet.BlockV2 (*class in haiku.nets*), 98
 ResNet101 (*class in haiku.nets*), 101
 ResNet152 (*class in haiku.nets*), 102
 ResNet18 (*class in haiku.nets*), 100
 ResNet200 (*class in haiku.nets*), 102
 ResNet34 (*class in haiku.nets*), 100
 ResNet50 (*class in haiku.nets*), 101
 reverse() (*haiku.nets.MLP* method), 96
 reverse_causal() (*in module haiku.pad*), 95
 RMSNorm (*class in haiku*), 69
 rng_reserve_size() (*in module haiku.experimental*), 121
 RNNCore (*class in haiku*), 73
 running_init() (*in module haiku*), 132
- ## S
- same() (*in module haiku.pad*), 95
 scan() (*in module haiku*), 107
 SeparableDepthwiseConv2D (*class in haiku*), 63
 Sequential (*class in haiku*), 53
 set() (*in module haiku.config*), 124
 set_policy() (*in module haiku.mixed_precision*), 111
 set_state() (*in module haiku*), 20
 SetterContext (*class in haiku*), 23
 shape (*haiku.experimental.ArraySpec* attribute), 116
 SNParamsTree (*class in haiku*), 71
 SpectralNorm (*class in haiku*), 70
 state (*haiku.experimental.ModuleDetails* attribute), 118
 State (*in module haiku*), 30
 state_dict() (*haiku.Module* method), 18
 static_unroll() (*in module haiku*), 74
 SupportsCall (*class in haiku*), 32
 switch() (*in module haiku*), 107
- ## T
- tabulate() (*in module haiku.experimental*), 114
 target (*haiku.Deferred* property), 88
 to_dot() (*in module haiku*), 48
 to_haiku_dict() (*in module haiku.data_structures*), 128
 to_immutable_dict() (*in module haiku.data_structures*), 129
 to_module() (*in module haiku*), 18
 to_mutable_dict() (*in module haiku.data_structures*), 129
 transform() (*in module haiku*), 12
 transform_and_run() (*in module haiku.testing*), 131
 transform_with_state() (*in module haiku*), 13
 Transformed (*class in haiku*), 30
 TransformedWithState (*class in haiku*), 30
 transparent() (*in module haiku*), 47
 transparent_lift() (*in module haiku*), 37

`transparent_lift_with_state()` (*in module haiku*),
37
`traverse()` (*in module haiku.data_structures*), 129
`tree_bytes()` (*in module haiku.data_structures*), 129
`tree_size()` (*in module haiku.data_structures*), 130
`TruncatedNormal` (*class in haiku.initializers*), 91

U

`UniformScaling` (*class in haiku.initializers*), 93

V

`valid()` (*in module haiku.pad*), 95
`value_and_grad()` (*in module haiku*), 109
`VanillaRNN` (*class in haiku*), 76
`VarianceScaling` (*class in haiku.initializers*), 92
`VectorQuantizer` (*class in haiku.nets*), 103
`VectorQuantizerEMA` (*class in haiku.nets*), 104
`vmap()` (*in module haiku*), 110

W

`while_loop()` (*in module haiku*), 108
`with_empty_state()` (*in module haiku*), 16
`with_rng()` (*in module haiku*), 28
`without_apply_rng()` (*in module haiku*), 16
`without_state()` (*in module haiku*), 16