

Arrays & Devices

Make arrays, see your accelerator, move data onto it.

PURPOSE	COMMAND	RESULT
Array from a Python list.	<code>jnp.array([1., 2., 3.])</code>	
Range / evenly spaced.	<code>jnp.arange(5) jnp.linspace(0,1,3)</code>	
Convert a NumPy array.	<code>jnp.asarray(np_array)</code>	
List the devices.	<code>jax.devices() jax.default_backend()</code>	
Put data on a device.	<code>jax.device_put(x, devices()[0])</code>	
Wait for async compute.	<code>x.block_until_ready()</code>	

a JAX array is immutable and lives on the accelerator; the default float dtype is float32, not float64

NumPy API & Immutable Updates

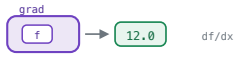





jnp mirrors numpy, but arrays never change in place.

PURPOSE	COMMAND	RESULT
Element-wise math (vectorized).	<code>jnp.exp(x) x * 2 + 1</code>	
Matrix multiply.	<code>a @ b jnp.dot(a, b)</code>	
Reduce over an axis.	<code>x.sum(axis=0) x.mean()</code>	
Indexing reads like numpy.	<code>m[1, 2] m[:, 0]</code>	
Functional "set" (returns copy).	<code>x.at[0].set(99)</code>	
Functional "add" at index.	<code>x.at[1].add(10)</code>	

arrays cannot be mutated: `x.at[idx].set(...)` / `.add(...)` returns a new array and leaves the original untouched

Automatic Differentiation


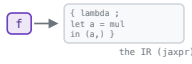
grad turns a function into its gradient function.

PURPOSE	COMMAND	RESULT
Gradient of a scalar function.	<code>grad(f)(2.0)</code>	
Value and gradient together.	<code>value_and_grad(loss)(w, x, y)</code>	
Differentiate chosen args.	<code>grad(loss, argnums=(0, 1))(w, b, x)</code>	
Higher-order derivative.	<code>grad(grad(f))(2.0)</code>	
Full Jacobian matrix.	<code>jax.jacobian(f)(x)</code>	
Hessian (curvature) matrix.	<code>jax.hessian(g)(x)</code>	

grad is a higher-order transform: give it a scalar function, get a new function computing the gradient

JIT Compilation






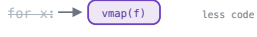
jit traces a function and compiles it with XLA.

PURPOSE	COMMAND	RESULT
Compile a function.	<code>jit(f)@jit</code>	
Trace once, then reuse.	<code>f(x); f(x)# 2nd call fast</code>	
Mark compile-time constants.	<code>partial(jit, static_argnums=(1,))</code>	
Inspect the traced program.	<code>jax.make_jaxpr(f)(x)</code>	
Time it correctly.	<code>f(x).block_until_ready()</code>	
Avoid: Python side effects.	<code>print()/ .item()</code> inside a jit'd fn	

jit traces once on shaped inputs and caches the XLA binary; the function must be pure, no Python if/for over array values

Auto-Vectorization (vmap)

Write the single-example function; vmap adds the batch axis.

PURPOSE	COMMAND	RESULT
Batch a per-example function.	<code>vmap(f)(batch)</code>	
Map over the leading axis.	<code>vmap(dot)(A, B)</code>	
Choose which axis to map.	<code>vmap(f, in_axes=(0, None))</code>	
Control the output axis.	<code>vmap(f, out_axes=1)</code>	
Compose with grad (per-sample).	<code>vmap(grad(f))(xs)</code>	
The payoff: no batch loop.	<code>for x in batch: vmap(f)</code>	

write a function for one example; in_axes/out_axes set the batch axis (None holds an arg fixed); composes with grad and jit

Random Numbers (Explicit Keys)


No global seed: you pass and split an explicit key.

PURPOSE	COMMAND	RESULT
Create a key from a seed.	<code>key =jax.random.key(0)</code>	
Split into independent keys.	<code>k1, k2 =jax.random.split(key)</code>	
Draw standard normals.	<code>jax.random.normal(key, (3,))</code>	
Uniform / integer draws.	<code>jax.random.uniform/ randint(key,...)</code>	
Per-step keys in a loop.	<code>jax.random.fold_in(key, step)</code>	
Same key, same draw.	<code>random.normal(key, (3,)) x2</code>	

the golden rule: never reuse a key; split it into fresh independent keys before each new draw

Control Flow & Loops


Inside jit, use lax primitives instead of Python branches.

PURPOSE	COMMAND	RESULT
Data-dependent branch.	<code>lax.cond(p, true_fn, false_fn, x)</code>	
Fixed-count loop.	<code>lax.fori_loop(0, 5, body, init)</code>	
Loop while a condition holds.	<code>lax.while_loop(cond, body, init)</code>	
Carry + collect (the RNN loop).	<code>lax.scan(body, init, xs)</code>	
Why not a Python for?	<code>forover a traced array -> unrolls</code>	
Stop a gradient flowing.	<code>jax.lax.stop_gradient(x)</code>	

inside a jit/grad trace, a Python if/for runs on the trace; data-dependent flow uses cond / fori_loop / while_loop / scan

Pytrees & Parameters

Nested dicts/lists of arrays are first-class; map over them.

PURPOSE	COMMAND	RESULT
A parameter pytree.	<code>params = {"w": ..., "b": ...}</code>	
Map a function over leaves.	<code>jax.tree.map(lambda x: x*2, params)</code>	
List just the leaves.	<code>jax.tree.leaves(params)</code>	
Flatten / unflatten.	<code>jax.tree.flatten(params)</code>	
One SGD step over a tree.	<code>jax.tree.map(lambda p,g: p - lr*g, params, grads)</code>	
Grad returns a matching tree.	<code>grad(loss)(params, batch)</code>	

a pytree is any nested dict/list/tuple of arrays; jax.tree.map applies over every leaf, preserving structure