

Keras

The essentials of Keras, the high-level way to build and train neural networks, from stacking layers to callbacks and transfer learning.

Build a Model







Three ways to assemble layers into a model.

PURPOSE	COMMAND	RESULT									
Linear stack of layers.	<code>model = keras.Sequential([...])</code>	 Sequential									
Declare the input shape.	<code>keras.Input(shape=(784,))</code>	 (None, 784) None = batch									
One dense hidden layer.	<code>keras.layers.Dense(128, "relu")</code>	 Dense(128) relu									
Functional API (a graph).	<code>outputs = keras.layers.Dense(10)(x)</code>	 Model(in, out)									
Subclass for full control.	<code>class Net(keras.Model): def call(self, x):</code>	 custom forward									
Inspect the architecture.	<code>model.summary()</code>	<table border="1" data-bbox="1136 989 1362 1041"> <thead> <tr> <th>Layer</th> <th>Output</th> <th>Param #</th> </tr> </thead> <tbody> <tr> <td>Dense</td> <td>(None, 128)</td> <td>100480</td> </tr> <tr> <td>Dense</td> <td>(None, 10)</td> <td>1290</td> </tr> </tbody> </table> Total params: 101,770	Layer	Output	Param #	Dense	(None, 128)	100480	Dense	(None, 10)	1290
Layer	Output	Param #									
Dense	(None, 128)	100480									
Dense	(None, 10)	1290									

Sequential is the quick stack; the functional API handles branches; subclassing gives a custom call

The Layer Catalog

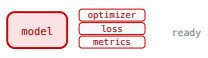



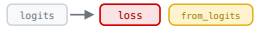
The building blocks you wire together.

PURPOSE	COMMAND	RESULT
Fully connected layer.	<code>keras.layers.Dense(64, "relu")</code>	 relu
2D convolution (images).	<code>keras.layers.Conv2D(32, 3, "relu")</code>	 32 maps
Downsample feature maps.	<code>keras.layers.MaxPooling2D(2)</code>	 max
Flatten / regularize.	<code>keras.layers.Flatten() Dropout(0.5)</code>	 Dropout
Sequence layers.	<code>Embedding(10000, 64) LSTM(64)</code>	 Embedding LSTM
Scale inputs in-graph.	<code>keras.layers.Rescaling(1./255)</code>	 preprocessing

Dense for tabular and heads; Conv2D/MaxPooling for images; Embedding/LSTM for sequences

Compile the Model



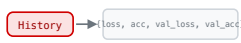



Choose the optimizer, loss, and metrics before training.

PURPOSE	COMMAND	RESULT
Wire up training.	<code>model.compile(optimizer, loss, metrics)</code>	
Pick an optimizer + LR.	<code>keras.optimizers.Adam(learning_rate=1e-3)</code>	
Multiclass loss (int labels).	<code>loss="sparse_categorical_crossentropy"</code>	
Binary / regression loss.	<code>"binary_crossentropy" "mse"</code>	
Track extra metrics.	<code>metrics=["accuracy", keras.metrics.AUC()]</code>	
Object form (fine control).	<code>keras.losses.SparseCategoricalCrossentropy(from_logits=True)</code>	

match the loss to the labels: `sparse_categorical` for int multiclass, `binary_crossentropy` for two-class, `mse` for regression

Train, Evaluate, Predict






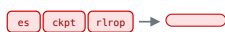
The fit / evaluate / predict loop over your data.

PURPOSE	COMMAND	RESULT
Train on arrays.	<code>model.fit(x, y, epochs=10, batch_size=32)</code>	
Hold out a validation set.	<code>model.fit(..., validation_split=0.2)</code>	
Capture the history.	<code>history = model.fit(...)</code>	
Score on test data.	<code>model.evaluate(x_test, y_test)</code>	
Predict new examples.	<code>preds = model.predict(x_new)</code>	
Feed a tf.data pipeline.	<code>model.fit(train_ds, epochs=10)</code>	

fit returns a History of per-epoch metrics; all three accept NumPy arrays or streamed tf.data datasets

Callbacks




Hooks that watch and steer training between epochs.

PURPOSE	COMMAND	RESULT
Stop when val stops improving.	<code>EarlyStopping(patience=3, restore_best_weights=True)</code>	 rewind to best
Save the best model so far.	<code>ModelCheckpoint("best.keras", save_best_only=True)</code>	 worse epochs skip
Drop LR on a plateau.	<code>ReduceLROnPlateau(factor=0.5, patience=2)</code>	 lr step down
Log to TensorBoard.	<code>TensorBoard(log_dir="logs")</code>	 dashboard
Append metrics to CSV.	<code>CSVLogger("training.csv")</code>	 training.csv
Pass the list to fit.	<code>model.fit(..., callbacks=[es, ckpt, rlrp])</code>	

collect callbacks in a list and hand it to `fit(..., callbacks=[...])`; they fire at batch and epoch edges

Save & Load


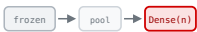
Persist a model, its weights, or its config to disk.

PURPOSE	COMMAND	RESULT
Save the whole model.	<code>model.save("model.keras")</code>	 arch + weights + compile
Load it back, ready to use.	<code>keras.models.load_model("model.keras")</code>	
Save weights only.	<code>model.save_weights("m.weights.h5")</code>	
Restore weights into a model.	<code>model.load_weights("m.weights.h5")</code>	 filled
Save architecture as JSON.	<code>json_str = model.to_json()</code>	 structure only
Export for serving.	<code>model.export("served_model")</code>	

.keras captures arch + weights + compile; save_weights stores only the numbers; export writes a SavedModel

Transfer Learning






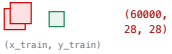
Reuse a pretrained network and fine-tune a small head.

PURPOSE	COMMAND	RESULT
Load a pretrained backbone.	<code>base = keras.applications.MobileNetV2(...)</code>	
Freeze the backbone.	<code>base.trainable = False</code>	
Add your own head.	<code>GlobalAveragePooling2D() +Dense</code>	
Preprocess to match the model.	<code>mobilenet_v2.preprocess_input(x)</code>	
Train just the head.	<code>model.compile(...); model.fit(epochs=5)</code>	
Unfreeze + fine-tune low LR.	<code>base.trainable = TrueAdam(1e-5)</code>	

freeze the backbone and train the head, then optionally unfreeze and fine-tune at a tiny learning rate

Backends & Config

One Keras 3 API, your choice of TensorFlow, JAX, or PyTorch.

PURPOSE	COMMAND	RESULT
Pick the backend (before import).	<code>os.environ["KERAS_BACKEND"] = "jax"</code>	
Confirm the active backend.	<code>keras.backend.backend()</code>	
Backend-agnostic tensor ops.	<code>keras.ops.matmul(a, b) ops.relu(x)</code>	
Make a run reproducible.	<code>keras.utils.set_random_seed(42)</code>	
Set default float precision.	<code>keras.config.set_floatx("float32")</code>	
Built-in datasets to start.	<code>keras.datasets.mnist.load_data()</code>	

set KERAS_BACKEND before importing keras; write once against keras.ops; reach for keras.datasets to experiment