

# Package ‘tensorflow’

May 8, 2026

**Type** Package

**Title** R Interface to 'TensorFlow'

**Version** 2.20.0

**Description** Interface to 'TensorFlow' <<https://www.tensorflow.org/>>, an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more 'CPUs' or 'GPUs' in a desktop, server, or mobile device with a single 'API'. 'TensorFlow' was originally developed by researchers and engineers working on the Google Brain Team within Google's Machine Intelligence research organization for the purposes of conducting machine learning and deep neural networks research, but the system is general enough to be applicable in a wide variety of other domains as well.

**License** Apache License 2.0

**URL** <https://github.com/rstudio/tensorflow>

**BugReports** <https://github.com/rstudio/tensorflow/issues>

**SystemRequirements** TensorFlow (<https://www.tensorflow.org/>)

**Encoding** UTF-8

**Depends** R (>= 3.6)

**Imports** config, processx, reticulate (>= 1.41.0), tfruns (>= 1.0),  
utils, yaml, grDevices, tfautograph (>= 0.3.1), rstudioapi (>= 0.7), lifecycle

**Suggests** testthat (>= 2.1.0), keras3, pillar, withr, callr

**RoxygenNote** 7.3.2

**NeedsCompilation** no

**Author** JJ Allaire [aut, cph],  
Tomasz Kalinowski [ctb, cph, cre],  
Daniel Falbel [ctb, cph],  
Dirk Eddelbuettel [ctb, cph],

Yuan Tang [aut, cph] (ORCID: <<https://orcid.org/0000-0001-5243-233X>>),  
 Nick Golding [ctb, cph],  
 Google Inc. [ctb, cph] (Examples and Tutorials),  
 Posit, PBC [cph, fnd]

**Maintainer** Tomasz Kalinowski <tomasz.kalinowski@posit.co>

**Repository** CRAN

**Date/Publication** 2025-08-22 15:10:08 UTC

## Contents

all_dims . . . . .	2
as_tensor . . . . .	3
evaluate . . . . .	4
install_tensorflow . . . . .	4
parse_arguments . . . . .	7
parse_flags . . . . .	8
set_random_seed . . . . .	9
shape . . . . .	10
tensorboard . . . . .	12
tensorflow . . . . .	13
tf . . . . .	14
tf_extract_opts . . . . .	15
tf_function . . . . .	16
tf_probability . . . . .	17
use_compat . . . . .	17
use_session_with_seed . . . . .	18
view_savedmodel . . . . .	19
[.tensorflow.tensor . . . . .	20
<b>Index</b>	<b>23</b>

---

all_dims	<i>All dims</i>
----------	-----------------

---

## Description

This function returns an object that can be used when subsetting tensors with `[]`. If you are familiar with python, this is equivalent to the python Ellipsis `...`, (not to be confused with `...` in R).

## Usage

```
all_dims()
```

**Examples**

```
## Not run:
# in python, if x is a numpy array or tensorflow tensor
x[..., i]
# the ellipsis means "expand to match number of dimension of x".
# to translate the above python expression to R, write:
x[all_dims(), i]

## End(Not run)
```

---

as\_tensor

*as\_tensor*


---

**Description**

Coerce objects to tensorflow tensors (potentially of a specific dtype or shape). The provided default methods will call `tf$convert_to_tensor`. Depending on arguments supplied it may also call some combination of

- `tf$saturate_cast` or `tf$cast`
- `tf$fill` or `tf$reshape`

**Usage**

```
as_tensor(x, dtype = NULL, ..., name = NULL)

## Default S3 method:
as_tensor(x, dtype = NULL, ..., shape = NULL, name = NULL)

## S3 method for class 'double'
as_tensor(x, dtype = NULL, ..., name = NULL)
```

**Arguments**

x	object to convert
dtype	NULL, a tensorflow dtype ( <code>tf\$int32</code> ), or something coercible to one (e.g. a string <code>"int32"</code> )
...	ignored
name	NULL or a string. Useful for debugging in graph mode, ignored while in eager mode.
shape	an integer vector, tensor, or <code>tf.TensorShape</code> . Can contain up to 1 unspecified dimension, encoded as a <code>-1</code> or <code>NA</code> . This will reshape <code>x</code> using row-major (C-style) semantics. It will prefer reshaping using non-graph operations if possible, but will otherwise invoke <code>tf\$reshape()</code> . If <code>x</code> is a scalar and the requested shape is fully defined or a tensor, the value of <code>x</code> will be recycled to fill a tensor of the requested shape (it will dispatch to <code>tf\$fill()</code> ).

**Value**

a tensorflow tensor

**Examples**

```
## Not run:
as_tensor(42, "int32")
as_tensor(as_tensor(42))

## End(Not run)
```

---

evaluate

*Evaluate a Model*

---

**Description**

Evaluate a model object. See implementations in the [keras3](#) package.

**Usage**

```
evaluate(object, ...)
```

**Arguments**

`object` An evaluatable R object.  
`...` Optional arguments passed on to implementing methods.

**Implementations**

- [keras3](#)

---

install\_tensorflow

*Install TensorFlow and its dependencies*

---

**Description**

This function installs TensorFlow into a persistent virtual environment. Beginning with reticulate version 1.41, in most circumstances, creating a persistent virtual environment by calling the `install_tensorflow()` function is no longer necessary, because reticulate automatically will resolve a python environment that satisfies all python requirements declared with `reticulate::py_require()`.

New code is recommended to call `py_require_tensorflow()` at the start of an R session to declare tensorflow requirements via `py_require()`. In a future package update this will by default be done in tensorflow's `.onLoad` hook.

The `py_require_tensorflow()` function that can dynamically modify the python requirements to enable usage of a GPU if one is available and usable by the R session.

The Python packages registered with `py_require()` by `py_require_tensorflow()`:

- On Linux: if a GPU is detected: "tensorflow[and-cuda]", otherwise, "tensorflow-cpu".
- On macOS: "tensorflow" is declared. The default package is not capable of using the GPU. To enable TensorFlow usage of the GPU, call `reticulate::py_require("tensorflow-metal")` before `reticulate` has initialized Python. Note that not all features of TensorFlow work correctly if `tensorflow-metal` is installed. There are known issues with random number generators like `tf$random$stateless_uniform()`, likely others as well.
- On Windows: "tensorflow" and "numpy<2" are declared. Note that TensorFlow GPU usage on Windows is no longer supported (Since TensorFlow 2.10). To use a GPU on windows, use TensorFlow via WSL. "numpy<2" is declared because at the time of this publishing, the pre-built binaries of tensorflow for Windows are not compatible with numpy>2.

`install_tensorflow()` creates a new virtual environment containing the tensorflow python package and it's direct dependencies. For creating a virtual environment with more complete set packages that includes additional optional dependencies, use [keras3::install\\_keras\(\)](#).

## Usage

```
install_tensorflow(
  method = c("auto", "virtualenv", "conda"),
  conda = "auto",
  version = "default",
  envname = "r-tensorflow",
  extra_packages = NULL,
  restart_session = TRUE,
  conda_python_version = NULL,
  ...,
  cuda = NULL,
  metal = FALSE,
  pip_ignore_installed = FALSE,
  new_env = identical(envname, "r-tensorflow"),
  python_version = NULL
)

py_require_tensorflow(use_gpu = NA)
```

## Arguments

method	Installation method. By default, "auto" automatically finds a method that will work in the local environment. Change the default to force a specific installation method. Note that the "virtualenv" method is not available on Windows.
conda	The path to a conda executable. Use "auto" to allow <code>reticulate</code> to automatically find an appropriate conda binary. See <b>Finding Conda</b> and <a href="#">conda_binary()</a> for more details.
version	TensorFlow version to install. Valid values include: <ul style="list-style-type: none"> <li>• "default" installs 2.20</li> <li>• "release" installs the latest release version of tensorflow (which may be incompatible with the current version of the R package)</li> </ul>

	<ul style="list-style-type: none"> <li>• A version specification like "2.4" or "2.4.0". Note that if the patch version is not supplied, the latest patch release is installed (e.g., "2.4" today installs version "2.4.2")</li> <li>• nightly for the latest available nightly build.</li> <li>• To any specification, you can append "-cpu" to install the cpu version only of the package (e.g., "2.4-cpu")</li> <li>• The full URL or path to an installer binary or python *.whl file.</li> </ul>
envname	The name, or full path, of the environment in which Python packages are to be installed. When NULL (the default), the active environment as set by the RETICULATE_PYTHON_ENV variable will be used; if that is unset, then the r-reticulate environment will be used.
extra_packages	Additional Python packages to install along with TensorFlow.
restart_session	Restart R session after installing (note this will only occur within RStudio).
conda_python_version	Passed to conda (only applicable if method = "conda")
...	other arguments passed to <code>reticulate::conda_install()</code> or <code>reticulate::virtualenv_install()</code> , depending on the method used.
cuda	logical TRUE or FALSE. If <code>install_tensorflow()</code> detects the platform is Linux, an Nvidia GPU is available, and the TensorFlow version is 2.14 (the default), it will install also install the required CUDA libraries through pip.
metal	Whether to install <code>tensorflow-metal</code> pip package on Arm Macs. This enables tensorflow to use the GPU. Pass a string to install a specific version like <code>"tensorflow-metal==0.7.*"</code> .
pip_ignore_installed	Whether pip should ignore installed python packages and reinstall all already installed python packages.
new_env	If TRUE, any existing Python virtual environment and/or conda environment specified by envname is deleted first.
python_version	Select the Python that will be used to create the virtualenv. Pass a string with version constraints like "3.8", or ">=3.9,<=3.11" or a file path to a python executable like <code>"/path/to/bin/python3"</code> . The supplied value is passed on to <code>reticulate::virtualenv_starter()</code> . Note that the Python version must be compatible with the requested TensorFlow version, documented here: <a href="https://www.tensorflow.org/install/pip#system-requirements">https://www.tensorflow.org/install/pip#system-requirements</a>
use_gpu	Only consulted if on Linux. It has no effect on macOS or Windows. If NA, the R package will attempt to detect a GPU. If a GPU is detected, then this is taken as TRUE, FALSE otherwise. If TRUE, then <code>tensorflow[and-cuda]</code> is declared, otherwise, <code>tensorflow-cpu</code> is declared via <code>py_require()</code> .

### Custom Installation

`install_tensorflow()` or `keras3::install_keras()` isn't required to use tensorflow with the package. If you manually configure a python environment with the required dependencies, you can tell R to use it by pointing `reticulate` at it, commonly by setting an environment variable:

```
Sys.setenv("RETICULATE_PYTHON" = "~/path/to/python-env/bin/python")
```

## Apple Silicon

Beginning with Tensorflow version 2.13, the default tensorflow package now works on Apple Silicon. See <https://developer.apple.com/metal/tensorflow-plugin/> for instructions on how to install older versions of Tensorflow on macOS. Please note that not all operations are supported on Arm Mac GPUs. You can work around the missing operations by pinning operations to CPU. For example:

```
x <- array(runif(64*64), c(1, 64, 64))
keras3::layer_random_rotation(x, .5) # Error:
# No registered 'RngReadAndSkip' OpKernel for 'GPU' devices
# Pin the operation to the CPU to avoid the error
with(tf$device("CPU"), keras3::layer_random_rotation(x, .5) ) # No Error
```

## Additional Packages

If you wish to add additional PyPI packages to your Keras / TensorFlow environment you can either specify the packages in the `extra_packages` argument of `install_tensorflow()` or `install_keras()`, or alternatively install them into an existing environment using the `reticulate::py_install()` function. Note that `install_keras()` includes a set of additional python packages by default, see `?keras3::install_keras` for details.

## See Also

- [keras3::install\\_keras\(\)](#)
- [https://tensorflow.rstudio.com/reference/tensorflow/install\\_tensorflow](https://tensorflow.rstudio.com/reference/tensorflow/install_tensorflow)

---

parse_arguments	<i>Parse Command Line Arguments</i>
-----------------	-------------------------------------

---

## Description

Parse command line arguments of the form `--key=value` and `--key value`. The values are assumed to be valid yaml and will be converted using `yaml::yaml.load()`.

## Usage

```
parse_arguments(arguments = NULL)
```

## Arguments

arguments	A vector of command line arguments. When NULL (the default), the command line arguments received by the current R process are used.
-----------	---

---

 parse\_flags

*Parse Configuration Flags for a TensorFlow Application*


---

## Description

Parse configuration flags for a TensorFlow application. Use this to parse and unify the configuration(s) specified through a `flags.yml` configuration file, alongside other arguments set through the command line.

## Usage

```
parse_flags(
  config = Sys.getenv("R_CONFIG_ACTIVE", unset = "default"),
  file = "flags.yml",
  arguments = commandArgs(TRUE)
)
```

## Arguments

<code>config</code>	The configuration to use. Defaults to the active configuration for the current environment (as specified by the <code>R_CONFIG_ACTIVE</code> environment variable), or default when unset.
<code>file</code>	The configuration file to read.
<code>arguments</code>	The command line arguments (as a character vector) to be parsed.

## Value

A named R list, mapping configuration keys to values.

## Examples

```
## Not run:
# examine an example configuration file provided by tensorflow
file <- system.file("examples/config/flags.yml", package = "tensorflow")
cat(readLines(file), sep = "\n")

# read the default configuration
FLAGS <- tensorflow::parse_flags("default", file = file)
str(FLAGS)

# read the alternate configuration: note that
# the default configuration is inherited, but
# we override the 'string' configuration here
FLAGS <- tensorflow::parse_flags("alternate", file = file)
str(FLAGS)

# override configuration values using command
# line arguments (normally, these would be
```

```
# passed in through the command line invocation
# used to start the process)
FLAGS <- tensorflow::parse_flags(
  "alternate",
  file = file,
  arguments = c("--foo=1")
)
str(FLAGS)

## End(Not run)
```

---

set_random_seed	<i>Set random seed for TensorFlow</i>
-----------------	---------------------------------------

---

## Description

Sets all random seeds needed to make TensorFlow code reproducible.

## Usage

```
set_random_seed(seed, disable_gpu = TRUE)
```

## Arguments

seed	A single value, interpreted as an integer
disable_gpu	TRUE to disable GPU execution (see <i>Parallelism</i> below).

## Details

This function should be used instead of `use_session_with_seed()` if you are using TensorFlow  $\geq 2.0$ , as the concept of session doesn't really make sense anymore.

This functions sets:

- The R random seed with `set.seed()`.
- The python and Numpy seeds via (`reticulate::py_set_seed()`).
- The TensorFlow seed with (`tf$random$set_seed()`)

It also optionally disables the GPU execution as this is a potential source of non-reproducibility.

---

shape *Create a tf.TensorShape object*

---

### Description

Create a `tf.TensorShape` object

### Usage

```
shape(..., dims = list(...))
```

### Arguments

`...` Tensor dimensions as integers or NULL for an unknown dimensions. NA and -1 are synonyms for NULL.

`dims` Tensor dimensions as a vector.

### See Also

[https://www.tensorflow.org/api\\_docs/python/tf/TensorShape](https://www.tensorflow.org/api_docs/python/tf/TensorShape)

### Examples

```
## Not run:

# --- construct ---
shape()      # tf.TensorShape()      # scalar
shape(NULL)  # tf.TensorShape([None]) # 1-D array of unknown length
shape(NA)    # tf.TensorShape([None]) # 1-D array of unknown length, NA is a synonym for NULL

shape(dims = NULL) # TensorShape(None) # Unknown rank, unknown size
shape(3, 4)        # TensorShape([3, 4]) # 2-D array (matrix) with 3 rows, 4 columns
shape(NA, 4)       # TensorShape([None, 4]) # 2-D array (matrix) with unknown rows, 4 columns
shape(dims = c(NA, 4)) # TensorShape([None, 4]) # same as above; bypass ... and pass dims directly

# --- inspect ---
length(shape(dims = NULL)) # NA_integer_
length(shape(1,2,3,NA))   # 4L

# ---convert ---
x <- shape(dims = list(3L, 5L))
as.list(x)      # list(3L, 5L)
as.integer(x)   # c(3L, 5L)
as.numeric(x)   # c(3, 5)
as.double(x)    # c(3, 5) # alias for as.numeric
as_tensor(x)    # tf.Tensor([3 5], shape=(2,), dtype=int32)

# convert partially undefined shapes
x <- shape(NA, 3)
```

```

as.list(x)      # list(NULL, 3L)
as.integer(x)  # c(NA, 3L)
as_tensor(x)   # tf.Tensor([-1 3], shape=(2,), dtype=int32) # unspecified dims default is -1

# as_tensor() converts undefined dimensions to -1, which is useful for
# tf functions that only accept tensors for shapes, e.g,
tf$reshape(tf$zeros(shape(8)),
           as_tensor(shape(NA, 4)))
# tf.Tensor([[0. 0. 0. 0.]
#           [0. 0. 0. 0.]], shape=(2, 4), dtype=float32)

# converting fully unknown shapes raises an error
try(as.list(shape(dims = NULL))) # ValueError: as_list() is not defined on an unknown TensorShape.
# test for rank first if this a concern:
as.list_or_null <- function(x) if(is.na(length(x))) NULL else as.list(x)
as.list_or_null(shape(dims = NULL))

# --- compare ---
# Fully known shapes return TRUE if and only if each element is equal
shape(3, 4) == shape(3, 4) # TRUE
shape(3, 4) == shape(4, 4) # FALSE

# two unknown dimensions are treated as equal
shape(NA, 4) == shape(NA, 4) # TRUE
shape(NA, 4) == shape(3, 4) # FALSE

# Two unknown shapes, return TRUE
shape(dims = NULL) == shape(dims = NULL) # TRUE

# Comparing an unknown shape to a partially or fully defined shape returns FALSE
shape(dims = NULL) == shape(NULL) # FALSE
shape(dims = NULL) == shape(4)   # FALSE

values of length greater than one supplied to `...` are automatically flattened
shape(1, c(2, 3), 4) # shape(1, 2, 3, 4)
shape(1, shape(2, 3), 4) # shape(1, 2, 3, 4)
shape(1, as_tensor(2, 3), 4) # shape(1, 2, 3, 4)

# --- extract or replace ---
# regular R-list semantics for `[`, `[[`, `[<-`, `[[<-`
x <- shape(1, 2, 3)
x[1]      # TensorShape([1])
x[[1]]    # 1L
x[2:3]    # TensorShape([2, 3])
x[-1]     # TensorShape([2, 3])

x[1] <- 11      ; x # TensorShape([11, 2, 3])
x[1] <- shape(11) ; x # TensorShape([11, 2, 3])
x[1] <- list(11) ; x # TensorShape([11, 2, 3])

x[[1]] <- 22      ; x # TensorShape([22, 2, 3])

```

```

x[1:2] <- c(NA, 99) ; x # TensorShape([None, 99, 3])
x[1:2] <- shape(33, 44) ; x # TensorShape([33, 44, 3])

# --- concatenate ---
c(shape(1), shape(2, 3), shape(4, NA)) # TensorShape([1, 2, 3, 4, None])

# --- merge ---
merge(shape(NA, 2),
      shape(1, 2)) # TensorShape([1, 2])

try(merge(shape(2, 2),
          shape(1, 2))) # ValueError: Shapes (2, 2) and (1, 2) are not compatible

rm(x) # cleanup

## End(Not run)

```

---

 tensorboard

*TensorBoard Visualization Tool*


---

## Description

TensorBoard is a tool inspecting and understanding your TensorFlow runs and graphs.

## Usage

```

tensorboard(
  log_dir,
  action = c("start", "stop"),
  host = "127.0.0.1",
  port = "auto",
  launch_browser = getOption("tensorflow.tensorboard.browser", interactive()),
  reload_interval = 5,
  purge_orphaned_data = TRUE
)

```

## Arguments

log_dir	Directories to scan for training logs. If this is a named character vector then the specified names will be used as aliases within TensorBoard.
action	Specify whether to start or stop TensorBoard (TensorBoard will be stopped automatically when the R session from which it is launched is terminated).
host	Host for serving TensorBoard
port	Port for serving TensorBoard. If "auto" is specified (the default) then an unused port will be chosen automatically.

- `launch_browser` Open a web browser for TensorBoard after launching. Defaults to TRUE in interactive sessions. When running under RStudio uses an RStudio window by default (pass a function e.g. `utils::browseURL()` to open in an external browser). Use the `tensorflow.tensorboard.browser` option to establish a global default behavior.
- `reload_interval` How often the backend should load more data.
- `purge_orphaned_data` Whether to purge data that may have been orphaned due to TensorBoard restarts. Disabling `purge_orphaned_data` can be used to debug data disappearance.

### Details

When TensorBoard is passed a `logdir` at startup, it recursively walks the directory tree rooted at `logdir` looking for subdirectories that contain tfevents data. Every time it encounters such a subdirectory, it loads it as a new run, and the frontend will organize the data accordingly.

The TensorBoard process will be automatically destroyed when the R session in which it is launched exits. You can pass `action = "stop"` to manually terminate TensorBoard.

### Value

URL for browsing TensorBoard (invisibly).

---

tensorflow	<i>TensorFlow for R</i>
------------	-------------------------

---

### Description

**TensorFlow** is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API.

### Details

The **TensorFlow API** is composed of a set of Python modules that enable constructing and executing TensorFlow graphs. The `tensorflow` package provides access to the complete TensorFlow API from within R.

For additional documentation on the `tensorflow` package see <https://tensorflow.rstudio.com>

### Author(s)

**Maintainer:** Tomasz Kalinowski <tomasz.kalinowski@posit.co> [contributor, copyright holder]

Authors:

- JJ Allaire [copyright holder]

- Yuan Tang <terrytangyuan@gmail.com> ([ORCID](#)) [copyright holder]

Other contributors:

- Daniel Falbel <daniel@posit.co> [contributor, copyright holder]
- Dirk Eddelbuettel <edd@debian.org> [contributor, copyright holder]
- Nick Golding <nick.golding.research@gmail.com> [contributor, copyright holder]
- Google Inc. (Examples and Tutorials) [contributor, copyright holder]
- Posit, PBC [copyright holder, funder]

### See Also

Useful links:

- <https://github.com/rstudio/tensorflow>
- Report bugs at <https://github.com/rstudio/tensorflow/issues>

---

tf

*Main TensorFlow module*

---

### Description

Interface to main TensorFlow module. Provides access to top level classes and functions as well as sub-modules (e.g. `tf$nn`, `tf$contrib$learn`, etc.).

### Usage

tf

### Format

TensorFlow module

### Examples

```
## Not run:
library(tensorflow)

hello <- tf$constant('Hello, TensorFlow!')
zeros <- tf$Variable(tf$zeros(shape(1L)))

tf$print(hello)
tf$print(zeros)

## End(Not run)
```

---

tf_extract_opts	<i>Tensor extract options</i>
-----------------	-------------------------------

---

**Description**

Tensor extract options

**Usage**

```
tf_extract_opts(
  style = getOption("tensorflow.extract.style"),
  ...,
  one_based = getOption("tensorflow.extract.one_based", TRUE),
  inclusive_stop = getOption("tensorflow.extract.inclusive_stop", TRUE),
  disallow_out_of_bounds = getOption("tensorflow.extract.disallow_out_of_bounds", TRUE),
  warn_tensors_passed_asis = getOption("tensorflow.extract.warn_tensors_passed_asis",
    TRUE),
  warn_negatives_pythonic = getOption("tensorflow.extract.warn_negatives_pythonic", TRUE)
)
```

**Arguments**

style	one of NULL (the default) "R" or "python". If supplied, this overrides all other options. "python" is equivalent to all the other arguments being FALSE. "R" is equivalent to warn_tensors_passed_asis and warn_negatives_pythonic set to FALSE
...	ignored
one_based	TRUE or FALSE, if one-based indexing should be used
inclusive_stop	TRUE or FALSE, if slices like start:stop should be inclusive of stop
disallow_out_of_bounds	TRUE or FALSE, whether checks are performed on the slicing index to ensure it is within bounds.
warn_tensors_passed_asis	TRUE or FALSE, whether to emit a warning the first time a tensor is supplied to [] that tensors are passed as-is, with no R to python translation
warn_negatives_pythonic	TRUE or FALSE, whether to emit a warning the first time a negative number is supplied to [] about the non-standard (python-style) interpretation

**Value**

an object with class "tf\_extract\_opts", suitable for passing to []. tensorflow.tensor()

**Examples**

```
## Not run:
x <- tf$constant(1:10)

opts <- tf_extract_opts("R")
x[1, options = opts]

# or for more fine-grained control
opts <- tf_extract_opts(
  one_based = FALSE,
  warn_tensors_passed_asis = FALSE,
  warn_negatives_pythonic = FALSE
)
x[0:2, options = opts]

## End(Not run)
```

---

 tf\_function

*Creates a callable TensorFlow graph from an R function.*


---

**Description**

tf\_function constructs a callable that executes a TensorFlow graph created by tracing the TensorFlow operations in f. This allows the TensorFlow runtime to apply optimizations and exploit parallelism in the computation defined by f.

**Usage**

```
tf_function(f, input_signature = NULL, autograph = TRUE, ...)
```

**Arguments**

f	the function to be compiled
input_signature	A possibly nested sequence of tf\$TensorSpec objects specifying the shapes and dtypes of the tensors that will be supplied to this function. If NULL, a separate function is instantiated for each inferred input signature. If input_signature is specified, every input to f must be a tensor.
autograph	TRUE or FALSE. If TRUE (the default), you can use tensors in R control flow expressions if, while, for and break and they will be traced into the tensorflow graph. A guide to getting started and additional details can be found: <a href="#">here</a>
...	additional arguments passed on to tf.function (vary based on Tensorflow version). See <a href="#">here</a> for details.

**Details**

A guide to getting started with [tf.function](#) can be found [here](#).

---

tf_probability	<i>TensorFlow Probability Module</i>
----------------	--------------------------------------

---

**Description**

TensorFlow Probability Module

**Usage**

```
tf_probability()
```

**Value**

Reference to [TensorFlow Probability](#) functions and classes

**Examples**

```
## Not run:  
library(tensorflow)  
## one time setup:  
# reticulate::py_install("tensorflow_probability")  
tfp <- tf_probability()  
tfp$distributions$Normal(loc = 0, scale = 1)  
  
## End(Not run)
```

---

use_compat	<i>Use Compatibility</i>
------------	--------------------------

---

**Description**

Enables TensorFlow to run under a different API version for compatibility with previous versions. For instance, this is useful to run TensorFlow 1.x code when using TensorFlow 2.x.

**Usage**

```
use_compat(version = c("v1", "v2"))
```

**Arguments**

version            The version to activate. Must be "v1" or "v2"

**Examples**

```
## Not run:
library(tensorflow)
use_compat("v1")

## End(Not run)
```

---

use\_session\_with\_seed *Use a session with a random seed*

---

**Description**

Set various random seeds required to ensure reproducible results. The provided seed value will establish a new random seed for R, Python, NumPy, and TensorFlow. GPU computations and CPU parallelism will also be disabled by default.

**Usage**

```
use_session_with_seed(
  seed,
  disable_gpu = TRUE,
  disable_parallel_cpu = TRUE,
  quiet = FALSE
)
```

**Arguments**

seed	A single value, interpreted as an integer
disable_gpu	TRUE to disable GPU execution (see <i>Parallelism</i> below).
disable_parallel_cpu	TRUE to disable CPU parallelism (see <i>Parallelism</i> below).
quiet	TRUE to suppress printing of messages.

**Details**

This function must be called at the very top of your script (i.e. immediately after `library(tensorflow)`, `library(keras)`, etc.). Any existing TensorFlow session is torn down via `tf$reset_default_graph()`.

This function takes all measures known to promote reproducible results from TensorFlow sessions, however it's possible that various individual TensorFlow features or dependent libraries escape its effects. If you encounter non-reproducible results please investigate the possible sources of the problem, contributions via pull request are very welcome!

Packages which need to be notified before and after the seed is set can register for the "tensorflow.on\_before\_use\_session" and "tensorflow.on\_use\_session" hooks (see [setHook\(\)](#) for additional details on hooks).

**Value**

TensorFlow session object, invisibly

**Parallelism**

By default the `use_session_with_seed()` function disables GPU and CPU parallelism, since both can result in non-deterministic execution patterns (see <https://stackoverflow.com/questions/42022950/>). You can optionally enable GPU or CPU parallelism by setting the `disable_gpu` and/or `disable_parallel_cpu` parameters to `FALSE`.

**Examples**

```
## Not run:  
library(tensorflow)  
use_session_with_seed(42)  
  
## End(Not run)
```

---

view_savedmodel	<i>View a Saved Model</i>
-----------------	---------------------------

---

**Description**

View a serialized model from disk.

**Usage**

```
view_savedmodel(model_dir)
```

**Arguments**

`model_dir`      The path to the exported model, as a string.

**Value**

URL for browsing TensorBoard (invisibly).

---

[.tensorflow.tensor    *Subset tensors with [*

---

## Description

Subset tensors with [

## Usage

```
## S3 method for class 'tensorflow.tensor'

x[
  ...,
  drop = TRUE,
  style = getOption("tensorflow.extract.style"),
  options = tf_extract_opts(style)
]
```

## Arguments

x	Tensorflow tensor
...	slicing specs. See examples and details.
drop	whether to drop scalar dimensions
style	One of "python" or "R".
options	An object returned by tf_extract_opts()

## Examples

```
## Not run:

x <- as_tensor(array(1:15, dim = c(3, 5)))
x
# by default, numerics supplied to [...] are interpreted R style
x[,1]    # first column
x[1:2,] # first two rows
x[,1, drop = FALSE] # 1 column matrix

# strided steps can be specified in R syntax or python syntax
x[, seq(1, 5, by = 2)]
x[, 1:5:2]
# if you are unfamiliar with python-style strided steps, see:
# https://numpy.org/doc/stable/reference/arrays.indexing.html#basic-slicing-and-indexing

# missing arguments for python syntax are valid, but they must by backticked
# or supplied as NULL
x[, `:::2`]
x[, NULL:NULL:2]
x[, `2:`]
```

```

# all_dims() expands to the shape of the tensor
# (equivalent to a python ellipsis `...`)
# (not to be confused with R dots `...`)
y <- as_tensor(array(1:(3^5), dim = c(3,3,3,3,3)))
all.equal(y[all_dims(), 1],
          y[, , , , 1])

# tf$newaxis are valid (equivalent to a NULL)
x[, , tf$newaxis]
x[, , NULL]

# negative numbers are always interpreted python style
# The first time a negative number is supplied to `[`, a warning is issued
# about the non-standard behavior.
x[-1,] # last row, with a warning
x[-1,] # the warning is only issued once

# specifying `style = 'python'` changes the following:
# + zero-based indexing is used
# + slice sequences in the form of `start:stop` do not include `stop`
#   in the returned value
# + out-of-bounds indices in a slice are valid

# The style argument can be supplied to individual calls of `[` or set
# as a global option

# example of zero based indexing
x[0, , style = 'python'] # first row
x[1, , style = 'python'] # second row

# example of slices with exclusive stop
options(tensorflow.extract.style = 'python')
x[, 0:1] # just the first column
x[, 0:2] # first and second column

# example of out-of-bounds index
x[, 0:10]
options(tensorflow.extract.style = NULL)

# slicing with tensors is valid too, but note, tensors are never
# translated and are always interpreted python-style.
# A warning is issued the first time a tensor is passed to `[`
x[, tf$constant(0L):tf$constant(2L)]
# just as in python, only scalar tensors are valid
# https://www.tensorflow.org/api_docs/python/tf/Tensor#__getitem__

# To silence the warnings about tensors being passed as-is and negative numbers
# being interpreted python-style, set
options(tensorflow.extract.style = 'R')

```

```
# clean up from examples
options(tensorflow.extract.style = NULL)

## End(Not run)
```

# Index

- \* **datasets**
  - tf, [14](#)
  - [.tensorflow.tensor, [20](#)
- all\_dims, [2](#)
- as\_tensor, [3](#)
  
- conda\_binary(), [5](#)
  
- evaluate, [4](#)
  
- install\_tensorflow, [4](#)
  
- keras3, [4](#)
- keras3::install\_keras(), [5](#), [7](#)
  
- parse\_arguments, [7](#)
- parse\_flags, [8](#)
- py\_require\_tensorflow
  - (install\_tensorflow), [4](#)
  
- reticulate::conda\_install(), [6](#)
- reticulate::py\_install(), [7](#)
- reticulate::py\_set\_seed(), [9](#)
- reticulate::virtualenv\_install(), [6](#)
  
- set.seed(), [9](#)
- set\_random\_seed, [9](#)
- setHook(), [18](#)
- shape, [10](#)
  
- tensorboard, [12](#)
- tensorflow, [13](#)
- tensorflow-package (tensorflow), [13](#)
- tf, [14](#)
- tf\_extract\_opts, [15](#)
- tf\_function, [16](#)
- tf\_probability, [17](#)
  
- use\_compat, [17](#)
- use\_session\_with\_seed, [18](#)
  
- use\_session\_with\_seed(), [9](#)
- utils::browseURL(), [13](#)
  
- view\_savedmodel, [19](#)
  
- yaml::yaml.load(), [7](#)