

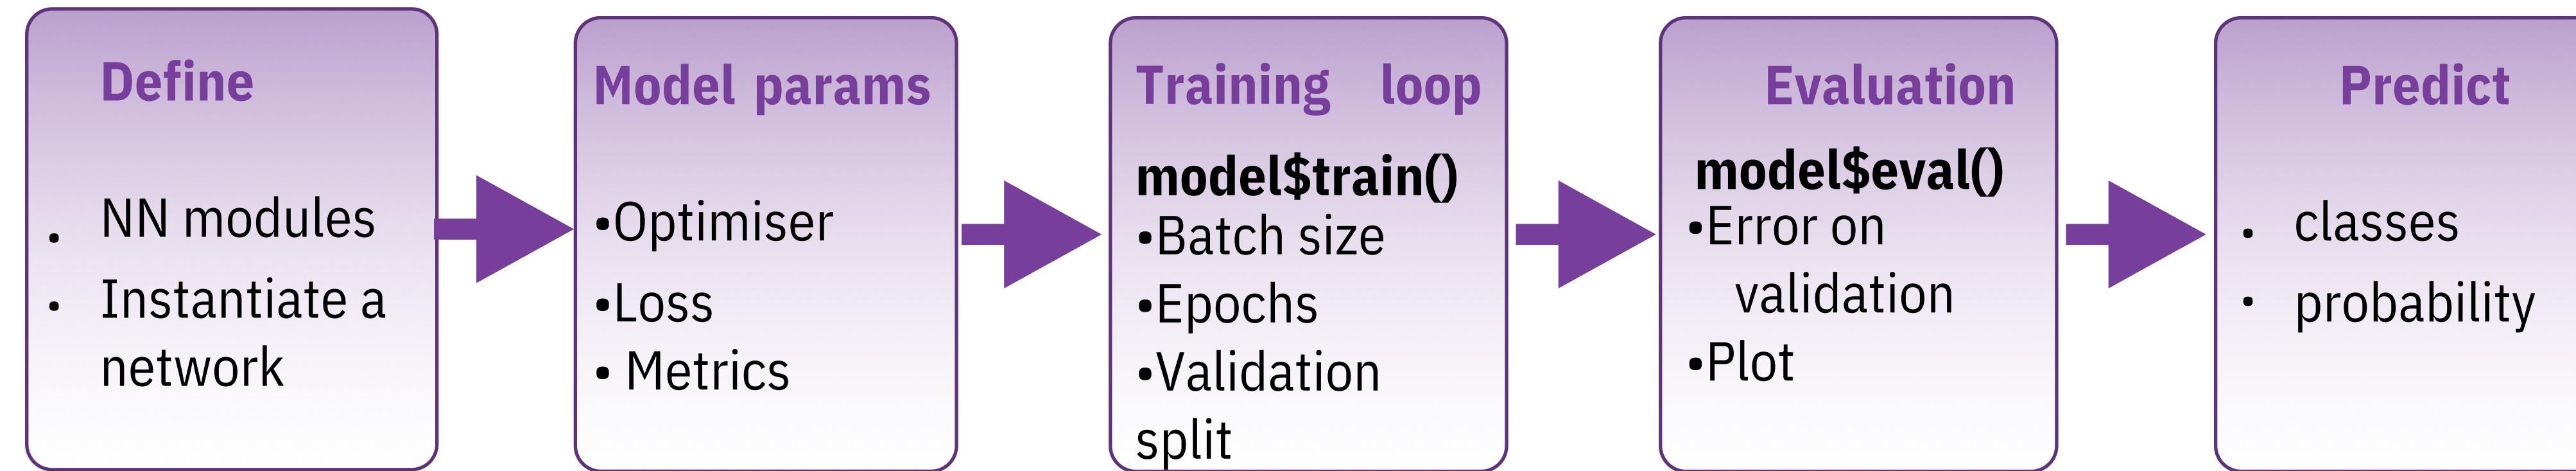
# Deep Learning with torch::

## Intro

{torch} is based on Pytorch, a framework popular among deep learning researchers.

{torch}'s GPU acceleration allows to implement fast machine learning algorithms using its convenient interface, as well as a vast range of use cases, not only for deep learning, according to its flexibility and its low level API.

It is part of an ecosystem of packages to interface with specific dataset like {torchaudio} for timeseries-like, {torchvision} for image-like, and {tabnet} for tabular data.



<https://torch.mlverse.org/>

<https://mlverse.shinyapps.io/torch-tour/>

## INSTALLATION

The torch R package uses the C++ libtorch library. You can install the prerequisites directly from R.

<https://torch.mlverse.org/docs/articles/installation.html>

```
install.packages("torch")
library(torch)
install_torch()
```

See `?install_torch` for GPU instructions

## Working with torch models

### DEFINE A NN MODULE

```
dense <- nn_module(  
  "no_bias_dense_layer",  
  initialize = function(in_f, out_f) {  
    self$w <- nn_parameter(torch_randn(in_f, out_f))  
  },  
  forward = function(x) {  
    torch_mm(x, self$w)  
  }  
)
```

Create a nn module names no\_bias\_dense\_layer

### ASSEMBLE MODULES INTO NETWORK

```
model <- dense(4, 3)  
Instantiate a network from a single module  
  
model <- nn_sequential(  
  dense(4,3), nn_relu(), nn_dropout(0.4),  
  dense(3,1), nn_sigmoid())  
Instantiate a sequential network with multiple layers
```

### MODEL FIT

```
model$train()  
Turns on gradient update  
  
with_enable_grad({  
  y_pred <- model(trainset)  
  loss <- (y_pred - y)$pow(2)$mean()  
  loss$backward()  
})  
Detailed training loop step (alternative)
```

### EVALUATE A MODEL

```
model$eval()  
or  
with_no_grad({  
  model(validationset)  
})  
Perform forward operation with no gradient update
```

### OPTIMIZATION

**optim\_sgd()**  
Stochastic gradient descent optimiser  
**optim\_adam()**  
ADAM optimiser

### CLASSIFICATION LOSS FUNCTION

**nn\_cross\_entropy\_loss()**  
**nn\_bce\_loss()**  
**nn\_bce\_with\_logits\_loss()**  
(Binary) cross-entropy losses  
**nn\_nll\_loss()**  
Negative log-likelihood loss  
**nn\_margin\_ranking\_loss()**  
**nn\_hinge\_embedding\_loss()**  
**nn\_multi\_margin\_loss()**  
**nn\_multilabel\_margin\_loss()**  
(Multiclass) (multi label) hinge losses

### REGRESSION LOSS FUNCTION

**nn\_l1\_loss()**  
L1 loss  
**nn\_mse\_loss()**  
MSE loss  
**nn\_ctc\_loss()**  
Connectionist Temporal Classification loss  
**nn\_cosine\_embedding\_loss()**  
Cosine embedding loss  
**nn\_kl\_div\_loss()**  
Kullback-Leibler divergence loss  
**nn\_poisson\_nll\_loss()**  
Poisson NLL loss

### OTHER MODEL OPERATIONS

**summary()** Print a summary of a torch model  
**torch\_save(); torch\_load()** Save/Load models to files  
**load\_state\_dict()**  
Load a model saved in python

## Neural-network layers

### CORE LAYERS

**nn\_linear()**  
Add a linear transformation NN layer to an input  
**nn\_bilinear()** to two inputs  
**nn\_sigmoid(), nn\_relu()**  
Apply an activation function to an output

### nn\_dropout()

**nn\_dropout2d()**  
**nn\_dropout3d()**  
Applies Dropout to the input

**nn\_batch\_norm1d()**  
**nn\_batch\_norm2d()**  
**nn\_batch\_norm3d()**

Applies batch normalisation to the weights

### CONVOLUTIONAL LAYERS

**convolution**  
**nn\_conv1d()** 1D, e.g. temporal

**nn\_conv\_transpose2d()**  
Transposed 2D (deconvolution)

**nn\_conv2d()** 2D, e.g. spatial convolution over images

**nn\_conv\_transpose3d()**  
Transposed 3D (deconvolution)  
**nn\_conv3d()** 3D, e.g. spatial convolution over volumes

**nff\_pad()**  
Zero-padding layer

### ACTIVATION LAYERS

**nn\_leaky\_relu()**  
Leaky version of a rectified linear unit

**nn\_relu6()**  
rectified linear unit clamped by 6

**nn\_rrelu()**  
Randomized leaky rectified linear unit

**nn\_elu(), nn\_selu()**  
Exponential linear unit, Scaled Exp lineal unit

### POOLING LAYERS

**nn\_max\_pool1d()**  
**nn\_max\_pool2d()**  
**nn\_max\_pool3d()**  
Maximum pooling for 1D to 3D

**nn\_avg\_pool1d()**  
**nn\_avg\_pool2d()**  
**nn\_avg\_pool3d()**  
Average pooling for 1D to 3D

**nn\_adaptive\_max\_pool1d()**  
**nn\_adaptive\_max\_pool2d()**  
**nn\_adaptive\_max\_pool3d()**  
Adaptive maximum pooling

**nn\_adaptive\_avg\_pool1d()**  
**nn\_adaptive\_avg\_pool2d()**  
**nn\_adaptive\_avg\_pool3d()**  
Adaptive average pooling

### RECURRENT LAYERS

**nn\_rnn()**  
Fully-connected RNN where the output is to be fed back to input

**nn\_gru()**  
Gated recurrent unit - Cho et al

**nn\_lstm()**  
Long-Short Term Memory unit - Hochreiter 1997  
CC BY SA Christophe Regouby • torch 0.7.0 • Updated: 2022-05

# Tensor manipulation

## TENSOR CREATION

```
tt <- torch_rand(4,3,2) uniform distrib.  
tt <- torch_rndn(4,3,2) unit normal distrib.  
tt <- torch_randint(1,7,c(4,3,2)) uniform integers within [1,7]  
Create a random values tensor with shape
```

```
tt <- torch_ones(4,3,2)
```

```
torch_ones_like(a)
```

Create a tensor full of 1 with given shape, or with the same shape as 'a'. Also **torch\_zeros**, **torch\_full**, **torch\_arange**...

```
tt$shape tt$ndim tt$dtype
```

```
[1] 4 3 2 [1] 3 torch_Float  
tt$requires_grad tt$device  
[1] FALSE torch_device(type='cpu')  
Get 't' tensor shape and attributes
```

```
tt$stride() jump needed to go from one
```

```
[1] 6 2 1 element to the next In each dimension
```

   
tt <- torch\_tensor(a,

```
dtype=torch_float(), device= "cuda")
```

Copy the R array 'a' into a tensor of float on the GPU

```
a <- as.matrix(tt$to(device="cpu"))
```

## TENSOR SLICING

```
tt[1:2, -2:-1, ]
```

Slice a 3D tensor

```
tt[5:N, -2:-1, ..]
```

Slice a 3D or more tensor, N for last

```
↓ -2 -1
```

```
1 tt[1:2, -2:-1, 1:1]
```

```
2 tt[1:2, -2:-1, 1, keep=TRUE]
```

NSlice a 3D and keep the unitary dim.

```
tt[1:2, -2:-1, 1]
```

Slice by default remove unitary dim.

   
tt[ tt > 3.1]

Boolean filtering (flattened result)

## TENSOR CONCATENATION

```
torch_stack()
```

Stack of tensors

```
torch_cat()
```

Assemble tensors

```
(, , ) (, , ) torch_split(2)
```

split tensor in sections of size 2

```
torch_split(c(1,3,1))
```

split tensor into explicit sizes

## TENSOR SHAPE OPERATIONS

```
tt$unsqueeze(1)
```

 Add a unitary dimension to tensor "tt" as first dimension

```
torch_squeeze(t,1)
```

 Remove first unitary dimension to tensor "tt"

```
tt$squeeze(1)
```

 Remove first unitary dimension to tensor "tt"

```
torch_reshape() $view()
```

Change the tensor shape, with copy or (tentatively) without

```
torch_flatten()
```

Flattens an input

```
torch_transpose()
```

switch dimension 1 with 2

```
torch_movedim(c(1,2), c(3,1,2))
```

move dim 1 to dim 3, dim 2 to 1, dim 3 to 2

```
torch_permute(c(3,1,2))
```

Only provide the target dimension order

```
torch_flip(1)
```

flip values along dim 1

```
torch_flip(2)
```

2

```
torch_flip(c(1,2))
```

both dims

## TENSOR VALUES OPERATIONS

```
+, -, *
```

Operations with two tensors

```
$pow(2), $log(), $exp(),
```

```
$abs(), $floor(), $round(), $cos(),
```

```
$fmod(3), $fmax(1), $fmin(3)
```

```
torch_clamp(tt, min=0.1, max=0.7)
```

Element-wise operations on a tensor

```
$eq(), $ge(), $le()
```

Element-wise comparison

```
$to(dtype = torch_long())
```

Mutate values type

```
$sum(dim=1), $mean(), $max()
```

Aggregation functions on a single tensor

```
$max()
```

nReports the input n times

```
torch_repeat_interleave()
```

nReports the input n times

## TRAINING AN IMAGE RECOGNIZER ON MNIST DATA

The "Hello, World!" of # input layer: use MNIST images

```
train_ds <- mnist_dataset(root = "~/.cache",  
download = TRUE,  
transform = torchvision::transform_to_tensor)
```

```
test_ds <- mnist_dataset(root = "~/.cache",  
train = FALSE,  
transform = torchvision::transform_to_tensor)
```

## Pre-trained models

Torch applications are deep learning models made available alongside pre-trained weights. These models can be used for prediction, feature extraction, and fine-tuning.

## NATIVE R MODELS

```
library(torchvision)  
resnet34 <- model_resnet34(pretrained=TRUE)  
Resnet image classification model  
resnet34_headless <- nn_prune_head(resnet34, 1)  
Remove top layer of a model
```

## IMPORTING FROM PYTORCH

{torchvisionlib} allows you to import pytorch model without recoding its nn module in R. This is done in two steps

```
1- instantiate the model in Python, script it, and save it.  
import torch  
import torchvision  
self$fc1 <- nn_linear(784, 128)  
self$fc2 <- nn_linear(128, 10)  
forward = function(x) {  
model = torchvision.models.segmentation.fcn_resnet50(pretrained = True)  
self$fc1() %>% nnf_relu() %>%  
model.eval()  
self$fc2() %>% nnf_log_softmax(dim = 1)  
scripted_model = torch.jit.script(model)  
torch.jit.save(scripted_model, "fcn_resnet50.pt")
```

```
2- load and use the model in R.  
model <- net()  
# define loss and optimizer
```

```
library(torchvisionlib)  
model <- torch::jit_load("fcn_resnet50.pt")  
# train (fit)
```

```
for (epoch in 1:10) {
```

```
train_losses <- c()
```

```
test_losses <- c()
```

```
for (b in enumerate(train_dl)) {  
optimizer$zero_grad()
```

```
output <- model(b[[1]]$to(device = device))
```

```
loss$backward()
```

```
optimizer$step()
```

```
train_losses <- c(train_losses, loss$item())
```

```
}
```

```
for (b in enumerate(test_dl)) {  
model$eval()
```

```
output <- model(b[[1]]$to(device = device))
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$train()
```

```
}
```

```
}
```

A callback is a set of functions to be applied at given stages of the training procedure. You can use them to get a view on internal states and statistics of the model during training.

```
callback <- function(model){  
model$b[[1]]$to(device = device)}
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$to(device = device)
```

```
loss <- nnf_nll_loss(output, b[[2]]$to(device = device))
```

```
test_losses <- c(test_losses, loss$item())
```

```
model$b[[1]]$
```