

# Package ‘cito’

October 6, 2023

**Type** Package

**Title** Building and Training Neural Networks

**Version** 1.0.2

**Description** Building and training custom neural networks in the typical R syntax. The 'torch' package is used for numerical calculations, which allows for training on CPU as well as on a graphics card.

**Encoding** UTF-8

**RoxygenNote** 7.2.3

**Depends** R (>= 3.5)

**Imports** coro, checkmate, torch, gridExtra, parabar, abind, progress, cli

**License** GPL (>= 3)

**Suggests** rmarkdown, testthat, plotly, ggraph, igraph, stats, ggplot2, knitr

**VignetteBuilder** knitr

**BugReports** <https://github.com/citoverse/cito/issues>

**URL** <https://citoverse.github.io/cito/>

**NeedsCompilation** no

**Author** Christian Amesöder [aut],  
Maximilian Pichler [aut, cre] (<<https://orcid.org/0000-0003-2252-8327>>),  
Florian Hartig [ctb] (<<https://orcid.org/0000-0002-6255-9059>>),  
Armin Schenk [ctb]

**Maintainer** Maximilian Pichler <[maximilian.pichler@biologie.uni-regensburg.de](mailto:maximilian.pichler@biologie.uni-regensburg.de)>

**Repository** CRAN

**Date/Publication** 2023-10-06 14:40:02 UTC

**R topics documented:**

ALE	2
analyze_training	4
cito	5
coef.citodnn	8
conditionalEffects	9
config_lr_scheduler	11
config_optimizer	13
continue_training	14
dnn	15
PDP	23
plot.citodnn	26
predict.citodnn	27
print.citodnn	28
print.conditionalEffects	29
print.summary.citodnn	29
residuals.citodnn	30
summary.citodnn	30
<b>Index</b>	<b>32</b>

ALE

*Accumulated Local Effect Plot (ALE)***Description**

Performs an ALE for one or more features.

**Usage**

```
ALE(
  model,
  variable = NULL,
  data = NULL,
  K = 10,
  ALE_type = c("equidistant", "quantile"),
  plot = TRUE,
  parallel = FALSE,
  ...
)
```

```
## S3 method for class 'citodnn'
```

```
ALE(
  model,
  variable = NULL,
  data = NULL,
  K = 10,
```

```

ALE_type = c("equidistant", "quantile"),
plot = TRUE,
parallel = FALSE,
...
)

## S3 method for class 'citodnnBootstrap'
ALE(
  model,
  variable = NULL,
  data = NULL,
  K = 10,
  ALE_type = c("equidistant", "quantile"),
  plot = TRUE,
  parallel = FALSE,
  ...
)

```

### Arguments

<code>model</code>	a model created by <code>dnn</code>
<code>variable</code>	variable as string for which the PDP should be done
<code>data</code>	data on which ALE is performed on, if NULL training data will be used.
<code>K</code>	number of neighborhoods original feature space gets divided into
<code>ALE_type</code>	method on how the feature space is divided into neighborhoods.
<code>plot</code>	plot ALE or not
<code>parallel</code>	parallelize over bootstrap models or not
<code>...</code>	arguments passed to <code>predict</code>

### Value

A list of plots made with 'ggplot2' consisting of an individual plot for each defined variable.

### Explanation

Accumulated Local Effect plots (ALE) quantify how the predictions change when the features change. They are similar to partial dependency plots but are more robust to feature collinearity.

### Mathematical details

If the defined variable is a numeric feature, the ALE is performed. Here, the non centered effect for feature  $j$  with  $k$  equally distant neighborhoods is defined as:

$$\hat{f}_{j,ALE}(x) = \sum_{k=1}^{k_j(x)} \frac{1}{n_j(k)} \sum_{i: x_j^{(i)} \in N_j(k)} \left[ \hat{f}(z_{k,j}, x_{\setminus j}^{(i)}) - \hat{f}(z_{k-1,j}, x_{\setminus j}^{(i)}) \right]$$

Where  $N_j(k)$  is the  $k$ -th neighborhood and  $n_j(k)$  is the number of observations in the  $k$ -th neighborhood.

The last part of the equation,  $\left[ \hat{f}(z_{k,j}, x_j^{(i)}) - \hat{f}(z_{k-1,j}, x_j^{(i)}) \right]$  represents the difference in model prediction when the value of feature  $j$  is exchanged with the upper and lower border of the current neighborhood.

### See Also

[PDP](#)

### Examples

```
if(torch::torch_is_installed()){
  library(cito)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris)

  ALE(nn.fit, variable = "Petal.Length")
}
```

---

analyze\_training

*Visualize training of Neural Network*

---

### Description

After training a model with cito, this function helps to analyze the training process and decide on best performing model. Creates a 'plotly' figure which allows to zoom in and out on training graph

### Usage

```
analyze_training(object)
```

### Arguments

object            a model created by [dnn](#)

### Details

The baseline loss is the most important reference. If the model was not able to achieve a better (lower) loss than the baseline (which is the loss for a intercept only model), the model probably did not converge. Possible reasons include an improper learning rate, too few epochs, or too much regularization. See the `?dnn` help or the vignette("B-Training\_neural\_networks").

### Value

a 'plotly' figure

## Examples

```
if(torch::torch_is_installed()){
  library(cito)
  set.seed(222)
  validation_set<- sample(c(1:nrow(datasets::iris)),25)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,],validation = 0.1)

  # show zoomable plot of training and validation losses
  analyze_training(nn.fit)

  # Use model on validation set
  predictions <- predict(nn.fit, iris[validation_set,])

  # Scatterplot
  plot(iris[validation_set,]$Sepal.Length,predictions)
}
```

---

cito

*'cito': Building and training neural networks*

---

## Description

'cito' simplifies the building and training of (deep) neural networks by relying on standard R syntax and familiar methods from statistical packages. Model creation and training can be done with a single line of code. Furthermore, all generic R methods such as print or plot can be used on the fitted model. At the same time, 'cito' is computationally efficient because it is based on the deep learning framework 'torch' (with optional GPU support). The 'torch' package is native to R, so no Python installation or other API is required for this package.

## Details

Cito is built around its main function `dnn`, which creates and trains a deep neural network. Various tools for analyzing the trained neural network are available.

## Installation

in order to install cito please follow these steps:

```
install.packages("cito")
library(torch)
install_torch(reinstall = TRUE)
library(cito)
```

### **cito functions and typical workflow**

- `dnn`: train deep neural network
- `analyze_training`: check for convergence by comparing training loss with baseline loss
- `continue_training`: continues training of an existing cito dnn model for additional epochs
- `summary.citodnn`: extract xAI metrics/effects to understand how predictions are made
- `PDP`: plot the partial dependency plot for a specific feature
- `ALE`: plot the accumulated local effect plot for a specific feature

Check out the vignettes for more details on training NN and how a typical workflow with 'cito' could look like.

### **Examples**

```
if(torch::torch_is_installed()){
  library(cito)

  # Example workflow in cito

  ## Build and train Network
  ### softmax is used for multi-class responses (e.g., Species)
  nn.fit<- dnn(Species~., data = datasets::iris, loss = "softmax")

  ## The training loss is below the baseline loss but at the end of the
  ## training the loss was still decreasing, so continue training for another 50
  ## epochs
  nn.fit <- continue_training(nn.fit, epochs = 50L)

  # Structure of Neural Network
  print(nn.fit)

  # Plot Neural Network
  plot(nn.fit)
  ## 4 Input nodes (first layer) because of 4 features
  ## 3 Output nodes (last layer) because of 3 response species (one node for each
  ## level in the response variable).
  ## The layers between the input and output layer are called hidden layers (two
  ## of them)

  ## We now want to understand how the predictions are made, what are the
  ## important features? The summary function automatically calculates feature
  ## importance (the interpretation is similar to an anova) and calculates
  ## average conditional effects that are similar to linear effects:
  summary(nn.fit)

  ## To visualize the effect (response-feature effect), we can use the ALE and
  ## PDP functions

  # Partial dependencies
  PDP(nn.fit, variable = "Petal.Length")
}
```

```

# Accumulated local effect plots
ALE(nn.fit, variable = "Petal.Length")

# Per se, it is difficult to get confidence intervals for our xAI metrics (or
# for the predictions). But we can use bootstrapping to obtain uncertainties
# for all cito outputs:
## Re-fit the neural network with bootstrapping
nn.fit<- dnn(Species~.,
            data = datasets::iris,
            loss = "softmax",
            epochs = 150L,
            verbose = FALSE,
            bootstrap = 20L)
## convergence can be tested via the analyze_training function
analyze_training(nn.fit)

## Summary for xAI metrics (can take some time):
summary(nn.fit)
## Now with standard errors and p-values
## Note: Take the p-values with a grain of salt! We do not know yet if they are
## correct (e.g. if you use regularization, they are likely conservative == too
## large)

## Predictions with bootstrapping:
dim(predict(nn.fit))
## The first dim corresponds to the bootstrapping, if you want the average
## predictions, you need to calculate the mean by your own:
apply(predict(nn.fit), 2:3, mean)

# Advanced: Custom loss functions and additional parameters
## Normal Likelihood with sd parameter:
custom_loss = function(true, pred) {
  logLik = torch::distr_normal(pred,
                              scale = torch::nnf_relu(scale)+
                              0.001)$log_prob(true)
  return(-logLik$mean())
}

nn.fit<- dnn(Sepal.Length~.,
            data = datasets::iris,
            loss = custom_loss,
            verbose = FALSE,
            custom_parameters = list(scale = 1.0)
)
nn.fit$parameter$scale

## Multivariate normal likelihood with parametrized covariance matrix
## Sigma = L*L^t + D
## Helper function to build covariance matrix

```

```

create_cov = function(LU, Diag) {
  return(torch::torch_matmul(LU, LU$t()) + torch::torch_diag(Diag+0.01))
}

custom_loss_MVN = function(true, pred) {
  Sigma = create_cov(SigmaPar, SigmaDiag)
  logLik = torch::distr_multivariate_normal(pred,
                                             covariance_matrix = Sigma)$
    log_prob(true)
  return(-logLik$mean())
}

nn.fit<- dnn(cbind(Sepal.Length, Sepal.Width, Petal.Length)~.,
             data = datasets::iris,
             lr = 0.01,
             verbose = FALSE,
             loss = custom_loss_MVN,
             custom_parameters =
               list(SigmaDiag = rep(1, 3),
                   SigmaPar = matrix(rnorm(6, sd = 0.001), 3, 2))
             )
as.matrix(create_cov(nn.fit$loss$parameter$SigmaPar,
                    nn.fit$loss$parameter$SigmaDiag))
}

```

---

coef.citodnn	<i>Returns list of parameters the neural network model currently has in use</i>
--------------	---

---

## Description

Returns list of parameters the neural network model currently has in use

## Usage

```

## S3 method for class 'citodnn'
coef(object, ...)

## S3 method for class 'citodnnBootstrap'
coef(object, ...)

```

## Arguments

object	a model created by <code>dnn</code>
...	nothing implemented yet



**Value**

list of weights of neural network

**Examples**

```
if(torch::torch_is_installed()){
  library(cito)

  set.seed(222)
  validation_set<- sample(c(1:nrow(datasets::iris)),25)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,])

  # Sturcture of Neural Network
  print(nn.fit)

  #analyze weights of Neural Network
  coef(nn.fit)
}
```

---

conditionalEffects      *Calculate average conditional effects*

---

**Description**

Average conditional effects calculate the local derivatives for each observation for each feature. They are similar to marginal effects. And the average of these conditional effects is an approximation of linear effects (see Pichler and Hartig, 2023 for more details). You can use this function to either calculate main effects (on the diagonal, take a look at the example) or interaction effects (off-diagonals) between features.

To obtain uncertainties for these effects, enable the bootstrapping option in the `dnn(. .)` function (see example).

**Usage**

```
conditionalEffects(
  object,
  interactions = FALSE,
  epsilon = 0.1,
  device = c("cpu", "cuda", "mps"),
  indices = NULL,
  data = NULL,
  type = "response",
  ...
)
```

```

## S3 method for class 'citodnn'
conditionalEffects(
  object,
  interactions = FALSE,
  epsilon = 0.1,
  device = c("cpu", "cuda", "mps"),
  indices = NULL,
  data = NULL,
  type = "response",
  ...
)

## S3 method for class 'citodnnBootstrap'
conditionalEffects(
  object,
  interactions = FALSE,
  epsilon = 0.1,
  device = c("cpu", "cuda", "mps"),
  indices = NULL,
  data = NULL,
  type = "response",
  ...
)

```

### Arguments

object	object of class citodnn
interactions	calculate interactions or not (computationally expensive)
epsilon	difference used to calculate derivatives
device	which device
indices	of variables for which the ACE are calculated
data	data which is used to calculate the ACE
type	ACE on which scale (response or link)
...	additional arguments that are passed to the predict function

### Value

an S3 object of class "conditionalEffects" is returned. The list consists of the following attributes:

result	3-dimensional array with the raw results
mean	Matrix, average conditional effects
abs	Matrix, summed absolute conditional effects
sd	Matrix, standard deviation of the conditional effects

**Author(s)**

Maximilian Pichler

**References**

Pichler, M., & Hartig, F. (2023). Can predictive models be used for causal inference?. arXiv preprint arXiv:2306.10551.

**Examples**

```
if(torch::torch_is_installed()){
  library(cito)

  # Build and train Network
  nn.fit = dnn(Sepal.Length~., data = datasets::iris)

  # Calculate average conditional effects
  ACE = conditionalEffects(nn.fit)

  ## Main effects (categorical features are not supported)
  ACE

  ## With interaction effects:
  ACE = conditionalEffects(nn.fit, interactions = TRUE)
  ## The off diagonal elements are the interaction effects
  ACE[[1]]$mean
  ## ACE is a list, elements correspond to the number of response classes
  ## Sepal.length == 1 Response so we have only one
  ## list element in the ACE object

  # Re-train NN with bootstrapping to obtain standard errors
  nn.fit = dnn(Sepal.Length~., data = datasets::iris, bootstrap = 30L)
  ## The summary method calculates also the conditional effects, and if
  ## bootstrapping was used, it will also report standard errors and p-values:
  summary(nn.fit)

}
```

---

config\_lr\_scheduler    *Creation of customized learning rate scheduler objects*

---

**Description**

Helps create custom learning rate schedulers for [dnn](#).

**Usage**

```
config_lr_scheduler(
  type = c("lambda", "multiplicative", "reduce_on_plateau", "one_cycle", "step"),
  verbose = FALSE,
  ...
)
```

**Arguments**

type	String defining which type of scheduler should be used. See Details.
verbose	If TRUE, additional information about scheduler will be printed to console.
...	additional arguments to be passed to scheduler. See Details.

**Details**

different learning rate scheduler need different variables, these functions will tell you which variables can be set:

- lambda: [lr\\_lambda](#)
- multiplicative: [lr\\_multiplicative](#)
- reduce\_on\_plateau: [lr\\_reduce\\_on\\_plateau](#)
- one\_cycle: [lr\\_one\\_cycle](#)
- step: [lr\\_step](#)

**Value**

object of class `cito_lr_scheduler` to give to [dnn](#)

**Examples**

```
if(torch::torch_is_installed()){
  library(cito)

  # create learning rate scheduler object
  scheduler <- config_lr_scheduler(type = "step",
    step_size = 30,
    gamma = 0.15,
    verbose = TRUE)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris, lr_scheduler = scheduler)
}
```

---

config\_optimizer      *Creation of customized optimizer objects*

---

## Description

Helps you create custom optimizer for [dnn](#). It is recommended to set learning rate in [dnn](#).

## Usage

```
config_optimizer(  
  type = c("adam", "adadelat", "adagrad", "rmsprop", "rprop", "sgd"),  
  verbose = FALSE,  
  ...  
)
```

## Arguments

type	character string defining which optimizer should be used. See Details.
verbose	If TRUE, additional information about scheduler will be printed to console
...	additional arguments to be passed to optimizer. See Details.

## Details

different optimizer need different variables, this function will tell you how the variables are set. For more information see the corresponding functions:

- adam: [optim\\_adam](#)
- adadelat: [optim\\_adadelat](#)
- adagrad: [optim\\_adagrad](#)
- rmsprop: [optim\\_rmsprop](#)
- rprop: [optim\\_rprop](#)
- sgd: [optim\\_sgd](#)

## Value

object of class cito\_optim to give to [dnn](#)

## Examples

```
if(torch::torch_is_installed()){  
  library(cito)  
  
  # create optimizer object  
  opt <- config_optimizer(type = "adagrad",  
                          lr_decay = 1e-04,
```

```

weight_decay = 0.1,
verbose = TRUE)

# Build and train Network
nn.fit<- dnn(Sepal.Length~., data = datasets::iris, optimizer = opt)

}

```

---

continue_training	<i>Continues training of a model generated with <a href="#">dnn</a> for additional epochs.</i>
-------------------	--

---

### Description

If the training/validation loss is still decreasing at the end of the training, it is often a sign that the NN has not yet converged. You can use this function to continue training instead of re-training the entire model.

### Usage

```

continue_training(
  model,
  epochs = 32,
  data = NULL,
  device = "cpu",
  verbose = TRUE,
  changed_params = NULL,
  parallel = FALSE
)

## S3 method for class 'citodnn'
continue_training(
  model,
  epochs = 32,
  data = NULL,
  device = "cpu",
  verbose = TRUE,
  changed_params = NULL,
  parallel = FALSE
)

## S3 method for class 'citodnnBootstrap'
continue_training(
  model,
  epochs = 32,
  data = NULL,
  device = "cpu",

```

```

    verbose = TRUE,
    changed_params = NULL,
    parallel = FALSE
  )

```

### Arguments

model	a model created by <a href="#">dnn</a>
epochs	additional epochs the training should continue for
data	matrix or data.frame if not provided data from original training will be used
device	device on which network should be trained on, either "cpu" or "cuda"
verbose	print training and validation loss of epochs
changed_params	list of arguments to change compared to original training setup, see <a href="#">dnn</a> which parameter can be changed
parallel	train bootstrapped model in parallel

### Value

a model of class `citodnn` or `citodnnBootstrap` created by [dnn](#)

### Examples

```

if(torch::torch_is_installed()){
  library(cito)

  set.seed(222)
  validation_set<- sample(c(1:nrow(datasets::iris)),25)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,], epochs = 32)

  # continue training for another 32 epochs
  nn.fit<- continue_training(nn.fit,epochs = 32)

  # Use model on validation set
  predictions <- predict(nn.fit, iris[validation_set,])
}

```

---

dnn

*DNN*


---

### Description

fits a custom deep neural network using the Multilayer Perceptron architecture. `dnn()` supports the formula syntax and allows to customize the neural network to a maximal degree.

**Usage**

```
dnn(
  formula,
  data = NULL,
  loss = c("mse", "mae", "softmax", "cross-entropy", "gaussian", "binomial", "poisson"),
  hidden = c(50L, 50L),
  activation = c("relu", "leaky_relu", "tanh", "elu", "rrelu", "prelu", "softplus",
    "celu", "selu", "gelu", "relu6", "sigmoid", "softsign", "hardtanh", "tanhshrink",
    "softshrink", "hardshrink", "log_sigmoid"),
  validation = 0,
  bias = TRUE,
  lambda = 0,
  alpha = 0.5,
  dropout = 0,
  optimizer = c("sgd", "adam", "adadelta", "adagrad", "rmsprop", "rprop"),
  lr = 0.01,
  batchsize = 32L,
  shuffle = TRUE,
  epochs = 100,
  bootstrap = NULL,
  bootstrap_parallel = FALSE,
  plot = TRUE,
  verbose = TRUE,
  lr_scheduler = NULL,
  custom_parameters = NULL,
  device = c("cpu", "cuda", "mps"),
  early_stopping = FALSE
)
```

**Arguments**

formula	an object of class " <b>formula</b> ": a description of the model that should be fitted
data	matrix or data.frame with features/predictors and response variable
loss	loss after which network should be optimized. Can also be distribution from the stats package or own function, see details
hidden	hidden units in layers, length of hidden corresponds to number of layers
activation	activation functions, can be of length one, or a vector of different activation functions for each layer
validation	percentage of data set that should be taken as validation set (chosen randomly)
bias	whether use biases in the layers, can be of length one, or a vector (number of hidden layers + 1 (last layer)) of logicals for each layer.
lambda	strength of regularization: lambda penalty, $\lambda * (L1 + L2)$ (see alpha)
alpha	add L1/L2 regularization to training $(1 - \alpha) *  weights  + \alpha   weights  ^2$ will get added for each layer. Can be single integer between 0 and 1 or vector of alpha values if layers should be regularized differently.
dropout	dropout rate, probability of a node getting left out during training (see <a href="#">nn_dropout</a> )



optimizer	which optimizer used for training the network, for more adjustments to optimizer see <a href="#">config_optimizer</a>
lr	learning rate given to optimizer
batchsize	number of samples that are used to calculate one learning rate step
shuffle	if TRUE, data in each batch gets reshuffled every epoch
epochs	epochs the training goes on for
bootstrap	bootstrap neural network or not, numeric corresponds to number of bootstrap samples
bootstrap_parallel	parallelize (CPU) bootstrapping
plot	plot training loss
verbose	print training and validation loss of epochs
lr_scheduler	learning rate scheduler created with <a href="#">config_lr_scheduler</a>
custom_parameters	List of parameters/variables to be optimized. Can be used in a custom loss function. See Vignette for example.
device	device on which network should be trained on. mps correspond to M1/M2 GPU devices.
early_stopping	if set to integer, training will stop if loss has gotten higher for defined number of epochs in a row, will use validation loss is available.

### Value

an S3 object of class "cito.dnn" is returned. It is a list containing everything there is to know about the model and its training process. The list consists of the following attributes:

net	An object of class "nn_sequential" "nn_module", originates from the torch package and represents the core object of this workflow.
call	The original function call
loss	A list which contains relevant information for the target variable and the used loss function
data	Contains data used for training the model
weights	List of weights for each training epoch
use_model_epoch	Integer, which defines which model from which training epoch should be used for prediction. 1 = best model, 2 = last model
loaded_model_epoch	Integer, shows which model from which epoch is loaded currently into model\$net.
model_properties	A list of properties of the neural network, contains number of input nodes, number of output nodes, size of hidden layers, activation functions, whether bias is included and if dropout layers are included.

**training\_properties**

A list of all training parameters that were used the last time the model was trained. It consists of learning rate, information about an learning rate scheduler, information about the optimizer, number of epochs, whether early stopping was used, if plot was active, lambda and alpha for L1/L2 regularization, batchsize, shuffle, was the data set split into validation and training, which formula was used for training and at which epoch did the training stop.

**losses**

A data.frame containing training and validation losses of each epoch


**Loss functions / Likelihoods**

We support loss functions and likelihoods for different tasks:

Name	Explanation	Example / Task
mse	mean squared error	Regression, predicting continuous values
mae	mean absolute error	Regression, predicting continuous values
softmax	categorical cross entropy	Multi-class, species classification
cross-entropy	categorical cross entropy	Multi-class, species classification
gaussian	Normal likelihood	Regression, residual error is also estimated (similar to <code>stats::lm()</code> )
binomial	Binomial likelihood	Classification/Logistic regression, mortality
poisson	Poisson likelihood	Regression, count data, e.g. species abundances

**Training and convergence of neural networks**

Ensuring convergence can be tricky when training neural networks. Their training is sensitive to a combination of the learning rate (how much the weights are updated in each optimization step), the batch size (a random subset of the data is used in each optimization step), and the number of epochs (number of optimization steps). Typically, the learning rate should be decreased with the size of the neural networks (depth of the network and width of the hidden layers). We provide a baseline loss (intercept only model) that can give hints about an appropriate learning rate:



If the training loss of the model doesn't fall below the baseline loss, the learning rate is either too high or too low. If this happens, try higher and lower learning rates.

A common strategy is to try (manually) a few different learning rates to see if the learning rate is on the right scale.

See the troubleshooting vignette (`vignette("B-Training_neural_networks")`) for more help on training and debugging neural networks.


### Finding the right architecture

As with the learning rate, there is no definitive guide to choosing the right architecture for the right task. However, there are some general rules/recommendations: In general, wider, and deeper neural networks can improve generalization - but this is a double-edged sword because it also increases the risk of overfitting. So, if you increase the width and depth of the network, you should also add regularization (e.g., by increasing the lambda parameter, which corresponds to the regularization strength). Furthermore, in [Pichler & Hartig, 2023](#), we investigated the effects of the hyperparameters on the prediction performance as a function of the data size. For example, we found that the selu activation function outperforms relu for small data sizes (<100 observations).

We recommend starting with moderate sizes (like the defaults), and if the model doesn't generalize/converge, try larger networks along with a regularization that helps minimize the risk of overfitting (see vignette("B-Training\_neural\_networks") ).

### Overfitting

Overfitting means that the model fits the training data well, but generalizes poorly to new observations. We can use the validation argument to detect overfitting. If the validation loss starts to increase again at a certain point, it often means that the models are starting to overfit your training data:



### Solutions:

- Re-train with epochs = point where model started to overfit
- Early stopping, stop training when model starts to overfit, can be specified using the `early_stopping=...` argument
- Use regularization (dropout or elastic-net, see next section)

## Regularization

Elastic Net regularization combines the strengths of L1 (Lasso) and L2 (Ridge) regularization. It introduces a penalty term that encourages sparse weight values while maintaining overall weight shrinkage. By controlling the sparsity of the learned model, Elastic Net regularization helps avoid overfitting while allowing for meaningful feature selection. We advise using elastic net (e.g.  $\lambda = 0.001$  and  $\alpha = 0.2$ ).

Dropout regularization helps prevent overfitting by randomly disabling a portion of neurons during training. This technique encourages the network to learn more robust and generalized representations, as it prevents individual neurons from relying too heavily on specific input patterns. Dropout has been widely adopted as a simple yet effective regularization method in deep learning.

By utilizing these regularization methods in your neural network training with the cito package, you can improve generalization performance and enhance the network's ability to handle unseen data. These techniques act as valuable tools in mitigating overfitting and promoting more robust and reliable model performance.

## Uncertainty

We can use bootstrapping to generate uncertainties for all outputs. Bootstrapping can be enabled by setting `bootstrap = ...` to the number of bootstrap samples to be used. Note, however, that the computational cost can be excessive.

In some cases it may be worthwhile to parallelize bootstrapping, for example if you have a GPU and the neural network is small. Parallelization for bootstrapping can be enabled by setting the `bootstrap_parallel = ...` argument to the desired number of calls to run in parallel.

## Custom Optimizer and Learning Rate Schedulers

When training a network, you have the flexibility to customize the optimizer settings and learning rate scheduler to optimize the learning process. In the cito package, you can initialize these configurations using the `config_lr_scheduler` and `config_optimizer` functions.

`config_lr_scheduler` allows you to define a specific learning rate scheduler that controls how the learning rate changes over time during training. This is beneficial in scenarios where you want to adaptively adjust the learning rate to improve convergence or avoid getting stuck in local optima.

Similarly, the `config_optimizer` function enables you to specify the optimizer for your network. Different optimizers, such as stochastic gradient descent (SGD), Adam, or RMSprop, offer various strategies for updating the network's weights and biases during training. Choosing the right optimizer can significantly impact the training process and the final performance of your neural network.

## How neural networks work

In Multilayer Perceptron (MLP) networks, each neuron is connected to every neuron in the previous layer and every neuron in the subsequent layer. The value of each neuron is computed using a weighted sum of the outputs from the previous layer, followed by the application of an activation function. Specifically, the value of a neuron is calculated as the weighted sum of the outputs of the neurons in the previous layer, combined with a bias term. This sum is then passed through an activation function, which introduces non-linearity into the network. The calculated value of each neuron becomes the input for the neurons in the next layer, and the process continues until the

output layer is reached. The choice of activation function and the specific weight values determine the network's ability to learn and approximate complex relationships between inputs and outputs.

Therefore the value of each neuron can be calculated using:  $a(\sum_j w_j * a_j)$ . Where  $w_j$  is the weight and  $a_j$  is the value from neuron  $j$  to the current one.  $a()$  is the activation function, e.g.  $relu(x) = max(0, x)$

### Training on graphic cards

If you have an NVIDIA CUDA-enabled device and have installed the CUDA toolkit version 11.3 and cuDNN 8.4, you can take advantage of GPU acceleration for training your neural networks. It is crucial to have these specific versions installed, as other versions may not be compatible. For detailed installation instructions and more information on utilizing GPUs for training, please refer to the [mlverse: 'torch' documentation](#).

Note: GPU training is optional, and the package can still be used for training on CPU even without CUDA and cuDNN installations.

### Author(s)

Christian Amesoeder, Maximilian Pichler

### See Also

[predict.citodnn](#), [plot.citodnn](#), [coef.citodnn](#), [print.citodnn](#), [summary.citodnn](#), [continue\\_training](#), [analyze\\_training](#), [PDP](#), [ALE](#),

### Examples

```
if(torch::torch_is_installed()){
  library(cito)

  # Example workflow in cito

  ## Build and train Network
  ### softmax is used for multi-class responses (e.g., Species)
  nn.fit<- dnn(Species~., data = datasets::iris, loss = "softmax")

  ## The training loss is below the baseline loss but at the end of the
  ## training the loss was still decreasing, so continue training for another 50
  ## epochs
  nn.fit <- continue_training(nn.fit, epochs = 50L)

  # Structure of Neural Network
  print(nn.fit)

  # Plot Neural Network
  plot(nn.fit)
  ## 4 Input nodes (first layer) because of 4 features
  ## 3 Output nodes (last layer) because of 3 response species (one node for each
  ## level in the response variable).
  ## The layers between the input and output layer are called hidden layers (two
```

```

## of them)

## We now want to understand how the predictions are made, what are the
## important features? The summary function automatically calculates feature
## importance (the interpretation is similar to an anova) and calculates
## average conditional effects that are similar to linear effects:
summary(nn.fit)

## To visualize the effect (response-feature effect), we can use the ALE and
## PDP functions

# Partial dependencies
PDP(nn.fit, variable = "Petal.Length")

# Accumulated local effect plots
ALE(nn.fit, variable = "Petal.Length")

# Per se, it is difficult to get confidence intervals for our xAI metrics (or
# for the predictions). But we can use bootstrapping to obtain uncertainties
# for all cito outputs:
## Re-fit the neural network with bootstrapping
nn.fit<- dnn(Species~.,
            data = datasets::iris,
            loss = "softmax",
            epochs = 150L,
            verbose = FALSE,
            bootstrap = 20L)
## convergence can be tested via the analyze_training function
analyze_training(nn.fit)

## Summary for xAI metrics (can take some time):
summary(nn.fit)
## Now with standard errors and p-values
## Note: Take the p-values with a grain of salt! We do not know yet if they are
## correct (e.g. if you use regularization, they are likely conservative == too
## large)

## Predictions with bootstrapping:
dim(predict(nn.fit))
## The first dim corresponds to the bootstrapping, if you want the average
## predictions, you need to calculate the mean by your own:
apply(predict(nn.fit), 2:3, mean)

# Advanced: Custom loss functions and additional parameters
## Normal Likelihood with sd parameter:
custom_loss = function(true, pred) {
  logLik = torch::distr_normal(pred,
                               scale = torch::nnf_relu(scale)+
                               0.001)$log_prob(true)
  return(-logLik$mean())
}

```

```

}

nn.fit<- dnn(Sepal.Length~.,
             data = datasets::iris,
             loss = custom_loss,
             verbose = FALSE,
             custom_parameters = list(scale = 1.0)
)
nn.fit$parameter$scale

## Multivariate normal likelihood with parametrized covariance matrix
## Sigma = L*L^t + D
## Helper function to build covariance matrix
create_cov = function(LU, Diag) {
  return(torch::torch_matmul(LU, LU$t()) + torch::torch_diag(Diag+0.01))
}

custom_loss_MVN = function(true, pred) {
  Sigma = create_cov(SigmaPar, SigmaDiag)
  logLik = torch::distr_multivariate_normal(pred,
                                             covariance_matrix = Sigma)$
    log_prob(true)
  return(-logLik$mean())
}

nn.fit<- dnn(cbind(Sepal.Length, Sepal.Width, Petal.Length)~.,
             data = datasets::iris,
             lr = 0.01,
             verbose = FALSE,
             loss = custom_loss_MVN,
             custom_parameters =
               list(SigmaDiag = rep(1, 3),
                   SigmaPar = matrix(rnorm(6, sd = 0.001), 3, 2))
)
as.matrix(create_cov(nn.fit$loss$parameter$SigmaPar,
                    nn.fit$loss$parameter$SigmaDiag))
}

```

**Description**

Calculates the Partial Dependency Plot for one feature, either numeric or categorical. Returns it as a plot.

**Usage**

```

PDP(
  model,
  variable = NULL,
  data = NULL,
  ice = FALSE,
  resolution.ice = 20,
  plot = TRUE,
  parallel = FALSE,
  ...
)

## S3 method for class 'citodnn'
PDP(
  model,
  variable = NULL,
  data = NULL,
  ice = FALSE,
  resolution.ice = 20,
  plot = TRUE,
  parallel = FALSE,
  ...
)

## S3 method for class 'citodnnBootstrap'
PDP(
  model,
  variable = NULL,
  data = NULL,
  ice = FALSE,
  resolution.ice = 20,
  plot = TRUE,
  parallel = FALSE,
  ...
)

```

**Arguments**

<code>model</code>	a model created by <a href="#">dnn</a>
<code>variable</code>	variable as string for which the PDP should be done. If none is supplied it is done for all variables.
<code>data</code>	specify new data PDP should be performed . If NULL, PDP is performed on the training data.
<code>ice</code>	Individual Conditional Dependence will be shown if TRUE
<code>resolution.ice</code>	resolution in which ice will be computed
<code>plot</code>	plot PDP or not
<code>parallel</code>	parallelize over bootstrap models or not



... arguments passed to [predict](#)

### Value

A list of plots made with 'ggplot2' consisting of an individual plot for each defined variable.

### Description

Performs a Partial Dependency Plot (PDP) estimation to analyze the relationship between a selected feature and the target variable.

The PDP function estimates the partial function  $\hat{f}_S$ :

$$\hat{f}_S(x_S) = \frac{1}{n} \sum_{i=1}^n \hat{f}(x_S, x_C^{(i)})$$

with a Monte Carlo Estimation:

$\hat{f}_S(x_S) = \frac{1}{n} \sum_{i=1}^n \hat{f}(x_S, x_C^{(i)})$  using a Monte Carlo estimation method. It calculates the average prediction of the target variable for different values of the selected feature while keeping other features constant.

For categorical features, all data instances are used, and each instance is set to one level of the categorical feature. The average prediction per category is then calculated and visualized in a bar plot.

If the `ice` parameter is set to `TRUE`, the Individual Conditional Expectation (ICE) curves are also shown. These curves illustrate how each individual data sample reacts to changes in the feature value. Please note that this option is not available for categorical features. Unlike PDP, the ICE curves are computed using a value grid instead of utilizing every value of every data entry.

Note: The PDP analysis provides valuable insights into the relationship between a specific feature and the target variable, helping to understand the feature's impact on the model's predictions. If a categorical feature is analyzed, all data instances are used and set to each level. Then an average is calculated per category and put out in a bar plot.

If `ice` is set to `true` additional the individual conditional dependence will be shown and the original PDP will be colored yellow. These lines show, how each individual data sample reacts to changes in the feature. This option is not available for categorical features. Unlike PDP the ICE curves are computed with a value grid instead of utilizing every value of every data entry.

### See Also

[ALE](#)

### Examples

```
if(torch::torch_is_installed()){
  library(cito)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris)

  PDP(nn.fit, variable = "Petal.Length")
}
```

---

plot.citodnn	<i>Creates graph plot which gives an overview of the network architecture.</i>
--------------	--

---

### Description

Creates graph plot which gives an overview of the network architecture.

### Usage

```
## S3 method for class 'citodnn'
plot(x, node_size = 1, scale_edges = FALSE, ...)

## S3 method for class 'citodnnBootstrap'
plot(x, node_size = 1, scale_edges = FALSE, which_model = 1, ...)
```

### Arguments

x	a model created by <a href="#">dnn</a>
node_size	size of node in plot
scale_edges	edge weight gets scaled according to other weights (layer specific)
...	no further functionality implemented yet
which_model	which model from the ensemble should be plotted

### Value

A plot made with 'ggraph' + 'igraph' that represents the neural network

### Examples

```
if(torch::torch_is_installed()){
  library(cito)

  set.seed(222)
  validation_set<- sample(c(1:nrow(datasets::iris)),25)

  # Build and train Network
  nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,])

  plot(nn.fit)
}
```

---

predict.citodnn      *Predict from a fitted dnn model*

---

## Description

Predict from a fitted dnn model

## Usage

```
## S3 method for class 'citodnn'
predict(
  object,
  newdata = NULL,
  type = c("link", "response", "class"),
  device = c("cpu", "cuda", "mps"),
  ...
)

## S3 method for class 'citodnnBootstrap'
predict(
  object,
  newdata = NULL,
  type = c("link", "response", "class"),
  device = c("cpu", "cuda", "mps"),
  ...
)
```

## Arguments

object	a model created by <code>dnn</code>
newdata	new data for predictions
type	which value should be calculated, either raw response, output of link function or predicted class (in case of classification)
device	device on which network should be trained on.
...	additional arguments

## Value

prediction matrix

## Examples

```
if(torch::torch_is_installed()){
  library(cito)
```

```

set.seed(222)
validation_set<- sample(c(1:nrow(datasets::iris)),25)

# Build and train Network
nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,])

# Use model on validation set
predictions <- predict(nn.fit, iris[validation_set,])
# Scatterplot
plot(iris[validation_set,]$Sepal.Length,predictions)
# MAE
mean(abs(predictions-iris[validation_set,]$Sepal.Length))
}

```

---

```
print.citodnn      Print class citodnn
```

---

## Description

Print class citodnn

## Usage

```

## S3 method for class 'citodnn'
print(x, ...)

## S3 method for class 'citodnnBootstrap'
print(x, ...)

```

## Arguments

```

x          a model created by dnn
...        additional arguments

```

## Value

prediction matrix  
original object x gets returned

## Examples

```

if(torch::torch_is_installed()){
  library(cito)

  set.seed(222)
  validation_set<- sample(c(1:nrow(datasets::iris)),25)

```

```
# Build and train Network
nn.fit<- dnn(Sepal.Length~., data = datasets::iris[-validation_set,])

# Sturcture of Neural Network
print(nn.fit)
}
```

---

```
print.conditionalEffects
      Print average conditional effects
```

---

### Description

Print average conditional effects

### Usage

```
## S3 method for class 'conditionalEffects'
print(x, ...)

## S3 method for class 'conditionalEffectsBootstrap'
print(x, ...)
```

### Arguments

x	print ACE calculated by <a href="#">conditionalEffects</a>
...	optional arguments for compatibility with the generic function, no function implemented

### Value

Matrix with average conditional effects

---

```
print.summary.citodnn Print method for class summary.citodnn
```

---

### Description

Print method for class summary.citodnn

### Usage

```
## S3 method for class 'summary.citodnn'
print(x, ...)

## S3 method for class 'summary.citodnnBootstrap'
print(x, ...)
```

**Arguments**

x                    a summary object created by [summary.citodnn](#)  
 ...                    additional arguments

**Value**

List with Matrices for importance, average CE, absolute sum of CE, and standard deviation of the CE

---

residuals.citodnn      *Extract Model Residuals*

---

**Description**

Returns residuals of training set.

**Usage**

```
## S3 method for class 'citodnn'
residuals(object, ...)
```

**Arguments**

object                a model created by [dnn](#)  
 ...                    no additional arguments implemented

**Value**

residuals of training set

---

summary.citodnn      *Summarize Neural Network of class citodnn*

---

**Description**

Performs a Feature Importance calculation based on Permutations

**Usage**

```
## S3 method for class 'citodnn'
summary(object, n_permute = NULL, device = NULL, ...)

## S3 method for class 'citodnnBootstrap'
summary(object, n_permute = NULL, device = NULL, ...)
```

**Arguments**

object	a model of class citodnn created by <a href="#">dnn</a>
n_permute	number of permutations performed. Default is $3 * \sqrt{n}$ , where n equals then number of samples in the training set
device	for calculating variable importance and conditional effects
...	additional arguments

**Details**

Performs the feature importance calculation as suggested by Fisher, Rudin, and Dominici (2018), and the mean and standard deviation of the average conditional Effects as suggested by Pichler & Hartig (2023).

Feature importances are in their interpretation similar to a ANOVA. Main and interaction effects are absorbed into the features. Also, feature importances are prone to collinearity between features, i.e. if two features are collinear, the importances might be overestimated.

Average conditional effects (ACE) are similar to marginal effects and approximate linear effects, i.e. their interpretation is similar to effects in a linear regression model.

The standard deviation of the ACE informs about the non-linearity of the feature effects. Higher values correlate with stronger non-linearities.

For each feature n permutation get done and original and permuted predictive mean squared error ( $e_{perm}$  &  $e_{orig}$ ) get evaluated with  $FI_j = e_{perm}/e_{orig}$ . Based on Mean Squared Error.

**Value**

summary.citodnn returns an object of class "summary.citodnn", a list with components

# Index

ALE, [2](#), [6](#), [21](#), [25](#)  
analyze\_training, [4](#), [6](#), [21](#)  
  
cito, [5](#)  
cito-package (cito), [5](#)  
coef.citodnn, [8](#), [21](#)  
coef.citodnnBootstrap (coef.citodnn), [8](#)  
conditionalEffects, [9](#), [29](#)  
config\_lr\_scheduler, [11](#), [17](#), [20](#)  
config\_optimizer, [13](#), [17](#), [20](#)  
continue\_training, [6](#), [14](#), [21](#)  
  
dnn, [3–6](#), [8](#), [11–15](#), [15](#), [24](#), [26–28](#), [30](#), [31](#)  
  
formula, [16](#)  
  
lr\_lambda, [12](#)  
lr\_multiplicative, [12](#)  
lr\_one\_cycle, [12](#)  
lr\_reduce\_on\_plateau, [12](#)  
lr\_step, [12](#)  
  
nn\_dropout, [16](#)  
  
optim\_adadelata, [13](#)  
optim\_adagrad, [13](#)  
optim\_adam, [13](#)  
optim\_rmsprop, [13](#)  
optim\_rprop, [13](#)  
optim\_sgd, [13](#)  
  
PDP, [4](#), [6](#), [21](#), [23](#)  
plot.citodnn, [21](#), [26](#)  
plot.citodnnBootstrap (plot.citodnn), [26](#)  
predict, [3](#), [25](#)  
predict.citodnn, [21](#), [27](#)  
predict.citodnnBootstrap  
    (predict.citodnn), [27](#)  
print.citodnn, [21](#), [28](#)  
print.citodnnBootstrap (print.citodnn),  
    [28](#)  
print.conditionalEffects, [29](#)  
print.conditionalEffectsBootstrap  
    (print.conditionalEffects), [29](#)  
print.summary.citodnn, [29](#)  
print.summary.citodnnBootstrap  
    (print.summary.citodnn), [29](#)  
  
residuals.citodnn, [30](#)  
  
summary.citodnn, [6](#), [21](#), [30](#), [30](#)  
summary.citodnnBootstrap  
    (summary.citodnn), [30](#)