

---

# **LossFunctions.jl Documentation**

***Release 0.0.3***

**Christof Stocker, Tom Breloff, Alex Williams**

December 02, 2016



<b>1</b>	<b>Getting Started</b>	<b>3</b>
1.1	Installation . . . . .	3
1.2	Hello World . . . . .	3
1.3	Overview . . . . .	4
1.4	Getting Help . . . . .	5
<b>2</b>	<b>Shared Interface</b>	<b>7</b>
2.1	Supervised Losses . . . . .	7
2.2	Distance-based Losses . . . . .	10
2.3	Margin-based Losses . . . . .	10
<b>3</b>	<b>Distance-based Losses</b>	<b>13</b>
3.1	LPDistLoss . . . . .	13
3.2	L1DistLoss . . . . .	13
3.3	L2DistLoss . . . . .	13
3.4	LogitDistLoss . . . . .	14
3.5	HuberLoss . . . . .	14
3.6	L1EpsilonInsLoss . . . . .	14
3.7	L2EpsilonInsLoss . . . . .	15
3.8	PeriodicLoss . . . . .	15
3.9	QuantileLoss . . . . .	15
<b>4</b>	<b>Margin-based Losses</b>	<b>17</b>
4.1	ZeroOneLoss . . . . .	17
4.2	PerceptronLoss . . . . .	17
4.3	L1HingeLoss . . . . .	17
4.4	L2HingeLoss . . . . .	18
4.5	LogitMarginLoss . . . . .	18
4.6	SmoothedL1HingeLoss . . . . .	18
4.7	ModifiedHuberLoss . . . . .	18
4.8	L2MarginLoss . . . . .	18
4.9	ExpLoss . . . . .	19
4.10	SigmoidLoss . . . . .	19
4.11	DWDMarginLoss . . . . .	19
<b>5</b>	<b>Indices and tables</b>	<b>21</b>
5.1	Acknowledgements . . . . .	21
5.2	LICENSE . . . . .	21

---



This package represents a community effort to centralize the definition and implementation of **loss functions** in Julia. As such, it is a part of the [JuliaML](#) ecosystem.

If this is the first time you consider using LossFunctions, make sure to check out the “Getting Started” section



---

## Getting Started

---

This section outlines the basic steps needed to start utilizing the `LossFunctions.jl` package. To that end we will provide a condensed overview of the package.

---

**Note:** In order to keep this overview concise, we will not discuss any background information or theory on the losses here in detail.

---

### 1.1 Installation

To install `LossFunctions.jl`, start up Julia and type the following code-snippet into the REPL. It makes use of the native Julia package manager.

```
Pkg.add("LossFunctions")
```

Additionally, for example if you encounter any sudden issues, or in the case you would like to contribute to the package, you can manually choose to be on the latest (untagged) version.

```
Pkg.checkout("LossFunctions")
```

### 1.2 Hello World

This package is registered in the Julia package ecosystem. Once installed the package can be imported just as any other Julia package.

```
using LossFunctions
```

The following code snippets show a simple scenario of how a *Loss* can be used to compute the element-wise values.

```
using LossFunctions

true_targets = [ 1, 0, -2]
pred_outputs = [0.5, 1, -1]

value(L2DistLoss(), true_targets, pred_outputs)
```

```
3-element Array{Float64,1}:
 0.25
```

```
1.0
1.0
```

Alternatively, one can also use the loss like a function to compute its `value()`.

```
myloss = L2DistLoss()
myloss(true_targets, pred_outputs) # same result as above
```

The function signatures of `value()` also apply to the derivatives.

```
deriv(L2DistLoss(), true_targets, pred_outputs)
```

```
3-element Array{Float64,1}:
-1.0
 2.0
 2.0
```

Additionally, we provide mutating versions of most functions.

```
buffer = zeros(3)
deriv!(buffer, L2DistLoss(), true_targets, pred_outputs)
```

If need be, one can also compute the `meanvalue()` or `sumvalue()` efficiently, without allocating a temporary array.

```
# or meanvalue
sumvalue(L2DistLoss(), true_targets, pred_outputs)
```

```
0.75
```

## 1.3 Overview

All the concrete loss “functions” that this package provides are defined as types and are subtypes of the abstract `Loss`. Typically the losses we work with in Machine Learning are bivariate functions of the true `target` and the predicted output of some prediction model. All losses that can be expressed this way are subtypes for `SupervisedLoss`. To compute the value of some `SupervisedLoss` we use the function `value()`.

```
value(L2DistLoss(), true_target, pred_output)
```

We can further divide the supervised losses into two useful sub-categories: `DistanceLoss` and `MarginLoss`.

### 1.3.1 Losses for Regression

Supervised losses that can be expressed as a univariate function of `output - target` are referred to as distance-based losses.

```
value(L2DistLoss(), difference)
```

Distance-based losses are typically utilized for regression problems. That said, there are also other losses that are useful for regression problems that don’t fall into this category, such as the `PeriodicLoss`.

---

**Note:** In the literature that this package is partially based on, the convention for the distance-based losses is `target - output` (see [STEINWART2008] p. 38). We chose to diverge from this definition because that would cause the the sign of the derivative to flip.

---



### 1.3.2 Losses for Classification

Margin-base losses are supervised losses where the values of the targets are restricted to be in  $\{-1, 1\}$ , and which can be expressed as a univariate function `output * target`.

```
value(L1HingeLoss(), agreement)
```

**Note:** Throughout the codebase we refer to the result of `output * target` as `agreement`. The discussion that lead to this convention can be found [issue #9](#)

Margin-based losses are usually used for binary classification. In contrast to other formalism, they do not natively provide probabilities as output.

**Note:** Even though distance-based losses and margin-based losses can be expressed in univariate form, we still provide the bivariate form of `value` for both.

## 1.4 Getting Help

To get help on specific functionality you can either look up the information here, or if you prefer you can make use of Julia's native doc-system. The following example shows how to get additional information on `L1HingeLoss` within Julia's REPL:


```
?L1HingeLoss
```

```
search: L1HingeLoss SmoothedL1HingeLoss
```

```
L1HingeLoss <: MarginLoss
```

The hinge loss linearly penalizes every prediction where the resulting agreement  $\leq 1$ . It is Lipschitz continuous and convex, but not strictly convex.

```
L(y, y) = max(0, 1 - yy)
```



If you find yourself stuck or have other questions concerning the package you can find us at [gitter](#) or the *Machine Learning* domain on [discourse.julialang.org](#)

- [Julia ML on Gitter](#)

- [Machine Learning on JuliaLang](#)

If you encounter a bug or would like to participate in the further development of this package come find us on Github.

- [JuliaML/LossFunctions.jl](#)

For details on a specific aspect, see the documentation outlined below.

---

## Shared Interface

---

All provided loss functions share a common interface with other losses, but to a varying degree.

This section will provide an overview of the shared functions are available to the different sub-class of an abstract `Loss`

### 2.1 Supervised Losses

#### **class** `SupervisedLoss`

Abstract subtype of `Loss`. A loss is considered **supervised**, if all the information needed to compute `value(loss, features, targets, outputs)` are contained in `targets` and `outputs`, and thus allows for the simplification `value(loss, targets, outputs)`.

#### 2.1.1 Computing the values

##### **value** (*loss, targets, outputs*)

Computes the value of the loss function for each observation-pair in `targets` and `outputs` individual and returns the result as an array of the same size as the parameters.

##### **Parameters**

- **loss** (*SupervisedLoss*) – An instance of the loss we are interested in.
- **targets** (*AbstractArray*) – The true targets that we would like your model to predict.
- **outputs** (*AbstractArray*) – The predicted outputs that your model produced.

**Returns** The values of the loss function for the elements in `targets` and `outputs`.

**Return type** `AbstractArray`

##### **sumvalue** (*loss, targets, outputs*)

Same as `value()`, but returns the **sum** of all values as a `Number` instead of all the individual values as `Array`.

**Return type** `Number`

##### **meanvalue** (*loss, targets, outputs*)

Same as `value()`, but returns the unweighted **mean** of all values as a single `Number` instead of all the individual values as `Array`.

**Return type** `Number`

**value!** (*buffer*, *loss*, *targets*, *outputs*)

Computes the values of the loss function for each observation-pair in *targets* and *outputs* individually and stores them in the preallocated *buffer*, which has to be the same size as the parameters.

**Parameters**

- **buffer** (*AbstractArray*) – Array to store the computed values in. Old values will be overwritten and lost.
- **loss** (*SupervisedLoss*) – An instance of the loss we are interested in.
- **targets** (*AbstractArray*) – The true targets that we would like your model to predict.
- **outputs** (*AbstractArray*) – The predicted outputs that your model produced.

**Returns** *buffer*

## 2.1.2 Computing the derivatives

**deriv** (*loss*, *targets*, *outputs*)

Computes the derivative of the loss function for each observation-pair in *targets* and *outputs* individually and returns the result as an array of the same size as the parameters.

**Parameters**

- **loss** (*SupervisedLoss*) – An instance of the loss we are interested in.
- **targets** (*AbstractArray*) – The true targets that we would like your model to predict.
- **outputs** (*AbstractArray*) – The predicted outputs that your model produced.

**Returns** The derivatives of the loss function for the elements in *targets* and *outputs*.

**Return type** *AbstractArray*

**sumderiv** (*loss*, *targets*, *outputs*)

Same as *deriv()*, but returns the **sum** of all derivatives as a single *Number*, instead of all the individual derivatives as *Array*.

**Return type** *Number*

**meanderiv** (*loss*, *targets*, *outputs*)

Same as *deriv()*, but returns the unweighted **mean** of all derivatives as a single *Number*, instead of all the individual derivatives as *Array*.

**Return type** *Number*

**deriv!** (*buffer*, *loss*, *targets*, *outputs*)

Computes the derivative of the loss function for each observation-pair in *targets* and *outputs* individually and stores them in the preallocated *buffer*, which has to be the same size as the parameters.

**Parameters**

- **buffer** (*AbstractArray*) – Array to store the computed derivatives in. Old values will be overwritten and lost.
- **loss** (*SupervisedLoss*) – An instance of the loss we are interested in.
- **targets** (*AbstractArray*) – The true targets that we would like your model to predict.
- **outputs** (*AbstractArray*) – The predicted outputs that your model produced.

**Returns** *buffer*

**value\_deriv**(*loss*, *targets*, *outputs*)

Returns the results of *value()* and *deriv()* as a tuple. In some cases this function can yield better performance, because the losses can make use of shared variable when computing the values.

### 2.1.3 Closures

In some circumstances it may be convenient to have the loss function or its derivative as a proper Julia function. Closures provide just that as the following examples demonstrate.

```
f = value_fun(L2DistLoss())
f(targets, outputs) # computes the value of L2DistLoss

d = deriv_fun(L2DistLoss())
d(targets, outputs) # computes the derivative of L2DistLoss
```

**value\_fun**(*loss*)

**deriv\_fun**(*loss*)

**deriv2\_fun**(*loss*)

**value\_deriv\_fun**(*loss*)

### 2.1.4 Querying loss properties

The losses implemented in this package provide a range of properties that can be queried by functions defined in *LearnBase.jl*.

**isminimizable**(*loss*)

**isconvex**(*loss*)

**isstrictlyconvex**(*loss*)

**isstronglyconvex**(*loss*)

**isdifferentiable**(*loss*[, *at*])

**istwicedifferentiable**(*loss*[, *at*])

**isnemitski**(*loss*)

**islipschitzcont**(*loss*)

**islocallylipschitzcont**(*loss*)

**isclipable**(*loss*)

**ismarginbased**(*loss*)

**isclasscalibrated**(*loss*)

**isdistancebased**(*loss*)

**issymmetric**(*loss*)

## 2.2 Distance-based Losses

### class `DistanceLoss`

Abstract subtype of `SupervisedLoss`. A supervised loss that can be simplified to  $L(\text{targets}, \text{outputs}) = L(\text{targets} - \text{outputs})$  is considered distance-based.

#### **value** (*loss*, *difference*)

Computes the value of the loss function for each observation in *difference* individually and returns the result as an array of the same size as the parameter.

##### Parameters

- **loss** (`DistanceLoss`) – An instance of the loss we are interested in.
- **difference** (`AbstractArray`) – The result of subtracting the true targets from the predicted outputs.

**Returns** The value of the loss function for the elements in *difference*.

**Return type** `AbstractArray`

#### **deriv** (*loss*, *difference*)

Computes the derivative of the loss function for each observation in *difference* individually and returns the result as an array of the same size as the parameter.

##### Parameters

- **loss** (`DistanceLoss`) – An instance of the loss we are interested in.
- **difference** (`AbstractArray`) – The result of subtracting the true targets from the predicted outputs.

**Returns** The derivatives of the loss function for the elements in *difference*.

**Return type** `AbstractArray`

#### **value\_deriv** (*loss*, *difference*)

Returns the results of `value()` and `deriv()` as a tuple. In some cases this function can yield better performance, because the losses can make use of shared variable when computing the values.

## 2.3 Margin-based Losses

### class `MarginLoss`

Abstract subtype of `SupervisedLoss`. A supervised loss, where the targets are in  $\{-1, 1\}$ , and which can be simplified to  $L(\text{targets}, \text{outputs}) = L(\text{targets} * \text{outputs})$  is considered margin-based.

#### **value** (*loss*, *agreement*)

Computes the value of the loss function for each observation in *agreement* individually and returns the result as an array of the same size as the parameter.

##### Parameters

- **loss** (`MarginLoss`) – An instance of the loss we are interested in.
- **agreement** (`AbstractArray`) – The result of multiplying the true targets with the predicted outputs.

**Returns** The value of the loss function for the elements in *agreement*.

**Return type** `AbstractArray`

**deriv** (*loss*, *agreement*)

Computes the derivative of the loss function for each observation in *agreement* individually and returns the result as an array of the same size as the parameter.

**Parameters**

- **loss** (*MarginLoss*) – An instance of the loss we are interested in.
- **agreement** (*AbstractArray*) – The result of multiplying the true targets with the predicted outputs.

**Returns** The derivatives of the loss function for the elements in *agreement*.

**Return type** *AbstractArray*

**value\_deriv** (*loss*, *agreement*)

Returns the results of *value()* and *deriv()* as a tuple. In some cases this function can yield better performance, because the losses can make use of shared variable when computing the values.





---

## Distance-based Losses

---

This section lists all the subtypes of *DistanceLoss* that are implemented in this package.

Fig. 3.1: Distance-based Losses (Regression)

### 3.1 LPDistLoss

Lossfunction	Derivative

**class LPDistLoss**

The  $p$ -th power absolute distance loss. It is Lipschitz continuous iff  $p = 1$ , convex if and only if  $p \geq 1$ , and strictly convex iff  $p > 1$ .

$$L(r) = |r|^p$$

### 3.2 L1DistLoss

Lossfunction	Derivative

**class L1DistLoss**

The absolute distance loss. Special case of the *LPDistLoss* with  $p=1$ . It is Lipschitz continuous and convex, but not strictly convex.

$$L(r) = |r|$$

### 3.3 L2DistLoss

Lossfunction	Derivative

**class L2DistLoss**

The least squares loss. Special case of the *LPDistLoss* with  $P=2$ . It is strictly convex.

$$L(r) = |r|^2$$

## 3.4 LogitDistLoss

Lossfunction	Derivative

**class LogitDistLoss**

The distance-based logistic loss for regression. It is strictly convex and Lipshitz continuous.

$$L(r) = -\ln \frac{4e^r}{(1+e^r)^2}$$

## 3.5 HuberLoss

Lossfunction	Derivative

**class HuberLoss**

$\alpha$

Loss function commonly used for robustness to outliers. For large values of  $\alpha$  it becomes close to the *L1DistLoss*, while for small values of  $\alpha$  it resembles the *L2DistLoss*. It is Lipshitz continuous and convex, but not strictly convex.

$$L(r) = \begin{cases} \frac{r^2}{2} & \text{if } |r| \leq \alpha \\ \alpha|r| - \frac{\alpha^2}{2} & \text{otherwise} \end{cases}$$

## 3.6 L1EpsilonInsLoss

Lossfunction	Derivative

**class L1EpsilonInsLoss**

The  $\epsilon$ -insensitive loss. Typically used in linear support vector regression. It ignores deviances smaller than  $\epsilon$ , but penalizes larger deviances linearly. It is Lipshitz continuous and convex, but not strictly convex.

$$L(r) = \max\{0, |r| - \epsilon\}$$

## 3.7 L2EpsilonInsLoss

Lossfunction	Derivative

**class** `L2EpsilonInsLoss`

The  $\epsilon$ -insensitive loss. Typically used in linear support vector regression. It ignores deviances smaller than  $\epsilon$ , but penalizes larger deviances quadratically. It is convex, but not strictly convex.

$$L(r) = \max\{0, |r| - \epsilon\}^2$$

## 3.8 PeriodicLoss

Lossfunction	Derivative

**class** `PeriodicLoss`

**c**

Measures distance on a circle of specified circumference  $c$ .

$$L(r) = 1 - \cos\left(\frac{2r\pi}{c}\right)$$

## 3.9 QuantileLoss

Lossfunction	Derivative

**class** `QuantileLoss`

**$\tau$**

The quantile loss, aka pinball loss. Typically used to estimate the conditional  $\tau$ -quantiles. It is convex, but not strictly convex. Furthermore it is Lipschitz continuous.

$$L(r) = \begin{cases} -(1 - \tau)r & \text{if } r < 0 \\ \tau r & \text{if } r \geq 0 \end{cases}$$



---

## Margin-based Losses

---

This section lists all the subtypes of *MarginLoss* that are implemented in this package.

Fig. 4.1: Margin-based Losses (Classification)

### 4.1 ZeroOneLoss

#### class **ZeroOneLoss**

The classical classification loss. It penalizes every missclassified observation with a loss of  $1$  while every correctly classified observation has a loss of  $0$ . It is not convex nor continuous and thus seldomly used directly. Instead one usually works with some classification-calibrated surrogate loss, such as one of those listed below.

$$L(a) = \begin{cases} 1 & \text{if } a < 0 \\ 0 & \text{if } a \geq 0 \end{cases}$$

### 4.2 PerceptronLoss

#### class **PerceptronLoss**

The perceptron loss linearly penalizes every prediction where the resulting agreement  $a \leq 0$ . It is Lipschitz continuous and convex, but not strictly convex.

$$L(a) = \max\{0, -a\}$$

### 4.3 L1HingeLoss

#### class **L1HingeLoss**

The hinge loss linearly penalizes every prediction where the resulting agreement  $a \leq 1$ . It is Lipschitz continuous and convex, but not strictly convex.

$$L(a) = \max\{0, 1 - a\}$$

## 4.4 L2HingeLoss

### class **L2HingeLoss**

The truncated least squares loss quadratically penalizes every prediction where the resulting agreement  $a \leq 1$ . It is locally Lipschitz continuous and convex, but not strictly convex.

$$L(a) = \max\{0, 1 - a\}^2$$

## 4.5 LogitMarginLoss

### class **LogitMarginLoss**

The margin version of the logistic loss. It is infinitely many times differentiable, strictly convex, and Lipschitz continuous.

$$L(a) = \ln(1 + e^{-a})$$

## 4.6 SmoothedL1HingeLoss

### class **SmoothedL1HingeLoss**

$\gamma$

As the name suggests a smoothed version of the L1 hinge loss. It is Lipschitz continuous and convex, but not strictly convex.

$$L(a) = \begin{cases} \frac{0.5}{\gamma} \cdot \max\{0, 1 - a\}^2 & \text{if } a \geq 1 - \gamma \\ 1 - \frac{\gamma}{2} - a & \text{otherwise} \end{cases}$$

## 4.7 ModifiedHuberLoss

### class **ModifiedHuberLoss**

A special (scaled) case of the *SmoothedL1HingeLoss* with  $\gamma = 4$ . It is Lipschitz continuous and convex, but not strictly convex.

$$L(a) = \begin{cases} \max\{0, 1 - a\}^2 & \text{if } a \geq -1 \\ -4a & \text{otherwise} \end{cases}$$

## 4.8 L2MarginLoss

### class **L2MarginLoss**

The margin-based least-squares loss for classification, which quadratically penalizes every prediction where  $a \neq 1$ . It is locally Lipschitz continuous and strongly convex.

$$L(a) = (1 - a)^2$$

## 4.9 ExpLoss

### class **ExpLoss**

The margin-based exponential Loss used for classification, which penalizes every prediction exponentially. It is infinitely many times differentiable, locally Lipschitz continuous and strictly convex, but not clipable.

$$L(a) = e^{-a}$$

## 4.10 SigmoidLoss

### class **SigmoidLoss**

The so called sigmoid loss is a continuous margin-base loss which penalizes every prediction with a loss within in the range (0,2). It is infinitely many times differentiable, Lipschitz continuous but nonconvex.

$$L(a) = 1 - \tanh(a)$$

## 4.11 DWDMarginLoss

### class **DWDMarginLoss**

**q**

The distance weighted discrimination margin loss. A differentiable generalization of the L1 hinge loss that is different than the *SmoothedL1HingeLoss*

$$L(a) = \begin{cases} 1 - a & \text{if } a \geq \frac{q}{q+1} \\ \frac{1}{a^q} \frac{q^q}{(q+1)^{q+1}} & \text{otherwise} \end{cases}$$





---

## Indices and tables

---

### 5.1 Acknowledgements

The basic design of this package is heavily modelled after the loss-related definitions in [\[STEINWART2008\]](#).

We would also like to mention that some early inspiration was drawn from [EmpiricalRisks.jl](#)

#### 5.1.1 References

### 5.2 LICENSE

The LossFunctions.jl package is licensed under the MIT “Expat” License  
see [LICENSE.md](#) in the Github repository.

- [genindex](#)
- [modindex](#)
- [search](#)



[STEINWART2008] Steinwart, Ingo, and Andreas Christmann. “Support vector machines”. Springer Science & Business Media, 2008.



## Symbols

$\alpha$  (HuberLoss attribute), 14  
 $\gamma$  (SmoothedL1HingeLoss attribute), 18  
 $\tau$  (QuantileLoss attribute), 15  
 (L1EpsilonInsLoss attribute), 14  
 (L2EpsilonInsLoss attribute), 15

## C

c (PeriodicLoss attribute), 15

## D

deriv  
     () (built-in function), 8  
 deriv() (built-in function), 8, 10  
 deriv2\_fun() (built-in function), 9  
 deriv\_fun() (built-in function), 9  
 DistanceLoss (built-in class), 10  
 DWDMarginLoss (built-in class), 19

## E

ExpLoss (built-in class), 19

## H

HuberLoss (built-in class), 14

## I

isclasscalibrated() (built-in function), 9  
 isclipable() (built-in function), 9  
 isconvex() (built-in function), 9  
 isdifferentiable() (built-in function), 9  
 isdistancebased() (built-in function), 9  
 islipschitzcont() (built-in function), 9  
 islocallylipschitzcont() (built-in function), 9  
 ismarginbased() (built-in function), 9  
 isminimizable() (built-in function), 9  
 isnemitski() (built-in function), 9  
 isstrictlyconvex() (built-in function), 9  
 isstronglyconvex() (built-in function), 9  
 issymmetric() (built-in function), 9  
 istwicedifferentiable() (built-in function), 9

## L

L1DistLoss (built-in class), 13  
 L1EpsilonInsLoss (built-in class), 14  
 L1HingeLoss (built-in class), 17  
 L2DistLoss (built-in class), 13  
 L2EpsilonInsLoss (built-in class), 15  
 L2HingeLoss (built-in class), 18  
 L2MarginLoss (built-in class), 18  
 LogitDistLoss (built-in class), 14  
 LogitMarginLoss (built-in class), 18  
 LPDistLoss (built-in class), 13

## M

MarginLoss (built-in class), 10  
 meanderiv() (built-in function), 8  
 meanvalue() (built-in function), 7  
 ModifiedHuberLoss (built-in class), 18

## P

PerceptronLoss (built-in class), 17  
 PeriodicLoss (built-in class), 15

## Q

q (DWDMarginLoss attribute), 19  
 QuantileLoss (built-in class), 15

## S

SigmoidLoss (built-in class), 19  
 SmoothedL1HingeLoss (built-in class), 18  
 sumderiv() (built-in function), 8  
 sumvalue() (built-in function), 7  
 SupervisedLoss (built-in class), 7

## V

value  
     () (built-in function), 7  
 value() (built-in function), 7, 10  
 value\_deriv() (built-in function), 8, 10, 11  
 value\_deriv\_fun() (built-in function), 9  
 value\_fun() (built-in function), 9

## Z

ZeroOneLoss (built-in class), [17](#)