
NumericExtensions.jl Documentation

Release 0.5

Dahua Lin

July 30, 2015

1	Package Overview	3
1.1	Motivating example	3
1.2	Main features	4
2	Functors	5
2.1	Motivating example	5
2.2	Functor types	6
2.3	Pre-defined functors	6
3	Mapping	9
3.1	General usage	9
3.2	Additional functions	9
3.3	Pre-defined mapping functions	10
3.4	Performance	10
4	Vector Broadcasting	13
5	Reduction	15
5.1	Basic reduction functions	15
5.2	Generic folding	17
5.3	Derived reduction functions	18
5.4	Weighted Sum	19
5.5	Performance	20
6	Vector Norms and Normalization	23
6.1	Vector Norm Evaluation	23
6.2	Normalization	23
7	Array Scanning	25
7.1	Generic Scanning	25
7.2	Common Scanning Functions	26

NumericExtensions.jl is a Julia package that provides high performance support of numerical computation. This package is an extension of the Julia Base – part of the material may be migrated into the Base in future.

NumericExtensions.jl provides a wide range of tools, which include:

- higher order functions for mapping, reduction, and map-reduce operation that takes typed functors to achieve performance comparable to hand-crafted loops.
- Functions that allow inplace updating and writing results to pre-allocated arrays for mapping, reduction, and map-reduce operations.
- Convenient functions for inplace vectorized computation.
- Vector broadcasting.
- Fast views for operating on contiguous blocks/slices.

Contents:

Package Overview

Julia provides a fantastic technical computing environment that allows you to write codes that are both performant and generic. However, as it is still at its early stage, some functions are not as performant as they can be and writing computational algorithms directly based on builtin functions may not give you the best performance. This package provides you with a variety of tools to address such issues.

1.1 Motivating example

To see how this package may help you, let's first consider a simple example, that is, to compute the sum of squared difference between two vectors. This can be easily done in Julia in one line as follows

```
r = sum(abs2(x - y))
```

Whereas this is simple, this expression involves some unnecessary operations that would lead to suboptimal performance: (1) it creates two temporary arrays respectively to store $x - y$ and $\text{abs}(x - y)$, (2) it completes the computation through three passes over the data – computing $x - y$, computing $\text{abs2}(x - y)$, and finally computing the sum. Julia provides a `mapreduce` function which allows you to complete the operation in a single pass without creating any temporaries:

```
r = mapreduce((x, y) -> abs2(x - y), +, x, y)
```

However, if you really run this you may probably find that this is even slower. The culprit here is that the anonymous function $(x, y) \rightarrow \text{abs2}(x - y)$ is not inlined; it will be resolved and called at each iteration. Therefore, to compute this efficiently, one has to write loops as below

```
s = 0.
for i = 1 : length(x)
    s += abs2(x[i] - y[i])
end
```

This is not too bad though, until you have more complex needs, e.g. computing this along each row/column of the matrix. Then writing the loops can become more involved, and it is more tricky to implement it in a cache-friendly way.

With this package, we can compute this efficiently without writing loops, as

```
r = sumfdiff(Abs2Fun(), x, y)

# to compute this along a specific dimension
r = sumfdiff(Abs2Fun(), x, y, dim)

# this package also provides sumsqdiff for this
```

```
r = sumabsdiff(x, y)
r = sumabsdiff(x, y, dim)
```

Here, `Abs2Fun` and `Add` are *typed functors* provided by this package, which, unlike normal functions, can still be properly inlined when passed into a higher order function (thus causing zero overhead). This package extends `map`, `foldl`, `foldr`, `sum`, and so on to accept typed functors and as well introduces additional higher order functions like `sumfdiff` and `scan` etc to simplify the usage in common cases.

Benchmark shows that writing in this way is over *9x faster* than `sum(abs2(x - y))`.

1.2 Main features

Main features of this package are highlighted below:

- Pre-defined functors that cover most typical mathematical computation;
- A easy way for user to define customized functors;
- Extended/specialized methods for `map`, `map!`, `foldl`, `foldr`, and `scan`. These methods are carefully optimized, which often result in $2x - 10x$ speed up;
- Additional functions such as `mapl!`, `sum!`, and `scan!` that allow inplace updating or writing results to preallocated arrays;
- Vector broadcasting computation (supporting both inplace updating and writing results to new arrays).
- Fast shared-memory views of arrays.

Many of the methods are extensions of base functions. Simply adding a statement using `NumericExtensions` is often enough for substantial performance improvement. Consider the following code snippet:

```
using NumericExtensions
x = rand(1000, 1000)
r = sum(x, 2)
```

Here, when adding the statement `using NumericExtensions` *transparently replace* the method provided in the `Base` module by the specialized method in *NumericExtensions*. As a consequence, the statement `r = sum(x, 2)` becomes *6x* faster. Using additional functions provided by this package can further improve the performance.

Functors

Passing functions as arguments are essential in writing generic algorithms. However, function arguments do not get inlined in Julia (at current version), usually resulting in suboptimal performance.

2.1 Motivating example

Passing functions as arguments are essential in writing generic algorithms. However, function arguments do not get inlined in Julia (at current version), usually resulting in suboptimal performance. Consider the following example:

```
plus(x, y) = x + y
map_plus(x, y) = map(plus, x, y)

a = rand(1000, 1000)
b = rand(1000, 1000)

# warming up and get map_plus compiled
a + b
map_plus(a, b)

# benchmark
@time for i in 1 : 10 map_plus(a, b) end # -- statement (1)
@time for i in 1 : 10 a + b end       # -- statement (2)
```

Run this script in your computer, you will find that statement (1) is over 20+ *times* slower than statement (2). The reason is that the function argument `plus` is resolved and called at each iteration of the inner loop within the `map` function.

This package addresses this issue through *type functors* (i.e. function-like objects of specific types) and a set of highly optimized higher level functions for mapping and reduction. The codes above can be rewritten as

```
using NumericExtensions

# benchmark
@time for i in 1 : 10 map(Add(), a, b) end # -- statement (1)
@time for i in 1 : 10 a + b end         # -- statement (2)
```

Here, using a typed functor `Add` statement (1) is 10% faster than statement (2) in my benchmark.

2.2 Functor types

Functor is the abstract base type for all functors, which are formally defined as below

```
abstract Functor{N} # N: the number of arguments

typealias UnaryFunctor Functor{1}
typealias BinaryFunctor Functor{2}
typealias TernaryFunctor Functor{3}
```

Below is an example that shows how to define a functor that computes the squared difference:

```
type SqrDiff <: Functor{2} end
NumericExtensions.evaluate(::SqrDiff, x, y) = abs2(x - y)
```

The package also provides macros `@functor1` and `@functor2`, respectively for defining unary and binary functors. For example,

```
# this defines a functor type MyAbs, which,
# when evaluated, invokes abs
@functor1 MyAbsFun abs

# this defines a functor type SqrDiff
sqrdiff(x, y) = abs2(x - y)
@functor2 SqrDiff sqrdiff
```

2.3 Pre-defined functors

NumericExtensions.jl has defined a series of functors as listed below:

- Arithmetic operators: Add, Subtract, Multiply, Divide, Negate, Pow, Modulo
- Comparison operators: Greater, GreaterEqual, Less, LessEqual, Equal, NotEqual
- Floating-point predicates: IsfiniteFun, IsinfFun, IsnanFun, IsequalFun
- Logical operators: Not, And, Or
- Bitwise operators: BitwiseNot, BitwiseAnd, BitwiseOr, BitwiseXor
- max and min: MaxFun, MinFun
- Rounding functors: FloorFun, CeilFun, RoundFun, TruncFun, IfloorFun, IceilFun, IroundFun, ItruncFun
- Algebraic functors: AbsFun, Abs2Fun, SqrFun, SqrtFun, CbrtFun, RcpFun, RsqrtFun, RcbtFun, HypotFun
- exp and log functors: ExpFun, Exp2Fun, Exp10Fun, LogFun, Log2Fun, Log10Fun, ExpmlFun, LoglpFun
- Trigonometric functors: SinFun, CosFun, TanFun, CotFun, CscFun, SecFun
- Inverse Trigonometric functors: AsinFun, AcosFun, AtanFun, Atan2Fun, AcotFun, AcscFun, AsecFun
- Hyperbolic functors: SinhFun, CoshFun, TanhFun, CothFun, CschFun, SechFun
- Inverse Hyperbolic functors: AsinhFun, AcoshFun, AtanhFun, AcothFun, AcschFun, AsechFun
- Error functors: ErfFun, ErfcFun, ErfInvFun, ErfcInvFun

- **Gamma functors:** GammaFun, LgammaFun, LfactFun, DigammaFun
- **Beta functors:** BetaFun, LbetaFun, EtaFun, ZetaFun
- **Airy functors:** AiryFun, AiryprimeFun, AiryaiFun, AiryaiprimeFun, AirybiFun, AirybiprimeFun
- **Bessel functors:** BesseljFun, Besselj0Fun, Besselj1Fun, BesseliFun, BesselkFun
- **Fused multiply and add:** FMA (i.e. $(a, b, c) \rightarrow a + b * c$)
- **Others:** LogitFun, LogisticFun, InvLogisticFun, XlogxFun, XlogyFun

Except for several functors that corresponding to operators, most functors are named using the capitalized version of the corresponding math function. Therefore, you don't have to look up this list to find the names. The collection of pre-defined functors will be extended in future. Please refer to `src/functors.jl` for the most updated list.

NumericExtensions.jl extends `map` and `map!` to accept functors for efficient element-wise mapping:

3.1 General usage

Synopsis:

Let `f1`, `f2`, and `f3` be respectively unary, binary, and ternary functors. Generic usage of `map` and `map!` is summarized as follows:

```
map(f1, x)
map(f2, x1, x2)
map(f3, x1, x2, x3)

map!(f1, dst, x)
map!(f2, dst, x1, x2)
map!(f3, dst, x1, x2, x3)
```

Here, `map` creates and returns the resultant array, while `map!` writes results to a pre-allocated `dst` and returns it. Each argument can be either an array or a scalar number. At least one argument should be an array, and all array arguments should have compatible sizes.

Examples:

```
map(AbsFun(), x)           # returns abs(x)
map(FMA(), x, y, z)       # returns x + y .* z
map!(Add(), dst, x, 2)    # writes x + 2 to dst
```

3.2 Additional functions

NumericExtensions.jl provides additional functions (`map1!`, `mapdiff`, and `mapdiff!`) to simplify common use:

Synopsis

`map1!` updates the first argument inplace with the results, `mapdiff` maps a functor to the difference between two arguments, and `mapdiff!` writes the results of `mapdiff` to a pre-allocated array.

```
map1!(f1, x1)              # x1 <-- f1(x1)
map1!(f2, x1, x2)         # x1 <-- f2(x1, x2)
map1!(f3, x1, x2, x3)     # x1 <-- f3(x1, x2, x3)
```

```
mapdiff(f1, x, y)           # returns f1(x - y)
mapdiff!(f1, dst, x, y)    # dst <-- f1(x - y)
```

Here, `x1` (*i.e.* the first argument to `map!`) must be an array, while `x2` and `x3` can be either an array or a number).

Note that `mapdiff` and `mapdiff!` uses an efficient implementation, which completes the computation in one-pass and never creates the intermediate array `x - y`.

Examples

```
map!(Mul(), x, 2)          # multiply x by 2 (inplace)
mapdiff(Abs2Fun(), x, y)   # compute squared differences between x and y
mapdiff(AbsFun(), x, 1)    # compute |x - 1|
```

3.3 Pre-defined mapping functions

Julia already provides vectorized function for most math computations. In this package, we additionally define several functions for vectorized inplace computation (based on `map!`), as follows

```
add!(x, y)                 # x <- x + y
subtract!(x, y)            # x <- x - y
multiply!(x, y)            # x <- x .* y
divide!(x, y)              # x <- x ./ y
negate!(x)                 # x <- -x
pow!(x, y)                 # x <- x .^ y

abs!(x)                    # x <- abs(x)
abs2!(x)                   # x <- abs2(x)
rcp!(x)                    # x <- 1 ./ x
sqrt!(x)                   # x <- sqrt(x)
exp!(x)                    # x <- exp(x)
log!(x)                    # x <- log(x)

floor!(x)                  # x <- floor(x)
ceil!(x)                   # x <- ceil(x)
round!(x)                  # x <- round(x)
trunc!(x)                  # x <- trunc(x)
```

In the codes above, `x` must be an array (*i.e.* an instance of `AbstractArray`), while `y` can be either an array or a scalar.

In addition, this package also define some useful functions using compound functors:

```
absdiff(x, y)              # abs(x - y)
sqrdiff(x, y)              # abs2(x - y)
fma(x, y, c)               # x + y .* c, where c can be array or scalar
fma!(x, y, c)              # x <- x + y .* c
```

3.4 Performance

For simple functions, such as `x + y` or `exp(x)`, the performance of the `map` version such as `map(Add(), x, y)` and `map(ExpFun(), x)` is comparable to the Julia counter part. However, `map` can accelerate computation considerably in a variety of cases:

- When the result storage has been allocated (e.g. in iterative updating algorithms) or you want inplace update, then `map!` or the pre-defined inplace computation function can be used to avoid unnecessary memory allocation/garbage collection, which can sometimes be the performance killer.
- When the inner copy contains two or multiple steps, `map` and `map!` can complete the computation in one-pass without creating intermediate arrays, usually resulting in about 2x or even more speed up. Benchmark shows that `absdiff(x, y)` and `sqrdiff(x, y)` are about 2.2x faster than `abs(x - y)` and `abs2(x - y)`.
- The script `test/benchmark_map.jl` runs a series of benchmarks to compare the performance `map` and the Julia vectorized expressions for a variety of computation.

Vector Broadcasting

Julia has very nice and performant functions for broadcasting: `broadcast` and `broadcast!`, which however does not work with functors. For customized computation, you have to pass in function argument, which would lead to severe performance degradation. *NumericExtensions.jl* provides `vbroadcast` and `vbroadcast!` to address this issue.

Synopsis

```
vbroadcast(f, x, y, dim) # apply a vector y to vectors of x along a specific dimension
vbroadcast!(f, dst, x, y, dim) # write results to dst
vbroadcast1!(f, x, y) # update x with the results
```

Here, `f` is a binary functor, and `x` and `y` are arrays such that `length(y) == size(x, dim)`.

Examples

```
vbroadcast(f, x, y, 1) # r[:,i] = f(x[:,i], y) for each i
vbroadcast(f, x, y, 2) # r[i,:] = f(x[i,:], y) for each i
vbroadcast(f, x, y, 3) # r[i,j,:] = f(x[i,j,:], y) for each i, j

vbroadcast1!(Add(), x, y, 1) # x[:,i] += y[:,i] for each i
vbroadcast1!(Mul(), x, y, 2) # x[i,:] .*= y[i,:] for each i
```

Difference from `broadcast`

Unlike `broadcast`, you have to specify the vector dimension for `vbroadcast`. The benefit is two-fold: (1) the overhead of figuring out broadcasting shape information is eliminated; (2) the shape of `y` can be flexible.

```
x = rand(5, 6)
y = rand(6)

vbroadcast(Add(), x, y, 2) # this adds y to each row of x
broadcast(+, x, reshape(y, 1, 6)) # with broadcast, you have to first reshape y into a row
```

Reduction

A key advantage of this package are highly optimized reduction and map-reduction functions, which sometimes lead to over 10x speed up.

5.1 Basic reduction functions

The package extends/specializes `sum`, `mean`, `max`, and `min`, not only providing substantially better performance, but also allowing reduction over function results. It also provides `sum!`, `mean!`, `max!`, and `min!`, which allow writing results to pre-allocated storage when performing reduction along specific dimensions.

The function `sum` and its variant forms:

```
sum(x)
sum(f1, x)           # compute sum of f1(x)
sum(f2, x1, x2)      # compute sum of f2(x1, x2)
sum(f3, x1, x2, x3)  # compute sum of f3(x1, x2, x3)

sum(x, dim)
sum(f1, x, dim)
sum(f2, x1, x2, dim)
sum(f3, x1, x2, x3, dim)

sum!(dst, x, dim)    # write results to dst
sum!(dst, f1, x1, dim)
sum!(dst, f2, x1, x2, dim)
sum!(dst, f3, x1, x2, x3, dim)

sumfdiff(f2, x, y)   # compute sum of f2(x - y)
sumfdiff(f2, x, y, dim)
sumfdiff!(dst, f2, x, y, dim)
```

The function `mean` and its variant forms:

```
mean(x)
mean(f1, x)          # compute mean of f1(x)
mean(f2, x1, x2)     # compute mean of f2(x1, x2)
mean(f3, x1, x2, x3) # compute mean of f3(x1, x2, x3)

mean(x, dim)
mean(f1, x, dim)
mean(f2, x1, x2, dim)
mean(f3, x1, x2, x3, dim)
```

```

mean!(dst, x, dim)      # write results to dst
mean!(dst, f1, x1, dim)
mean!(dst, f2, x1, x2, dim)
mean!(dst, f3, x1, x2, x3, dim)

meanfdiff(f2, x, y)     # compute mean of f2(x - y)
meanfdiff(f2, x, y, dim)
meanfdiff!(dst, f2, x, y, dim)

```

The function max and its variants:

```

max(x)
max(f1, x)              # compute maximum of f1(x)
max(f2, x1, x2)         # compute maximum of f2(x1, x2)
max(f3, x1, x2, x3)     # compute maximum of f3(x1, x2, x3)

max(x, (), dim)
max(f1, x, dim)
max(f2, x1, x2, dim)
max(f3, x1, x2, x3, dim)

max!(dst, x, dim)      # write results to dst
max!(dst, f1, x1, dim)
max!(dst, f2, x1, x2, dim)
max!(dst, f3, x1, x2, x3, dim)

maxfdiff(f2, x, y)     # compute maximum of f2(x - y)
maxfdiff(f2, x, y, dim)
maxfdiff!(dst, f2, x, y, dim)

```

The function min and its variants

```

min(x)
min(f1, x)              # compute minimum of f1(x)
min(f2, x1, x2)         # compute minimum of f2(x1, x2)
min(f3, x1, x2, x3)     # compute minimum of f3(x1, x2, x3)

min(x, (), dim)
min(f1, x, dim)
min(f2, x1, x2, dim)
min(f3, x1, x2, x3, dim)

min!(dst, x, dim)      # write results to dst
min!(dst, f1, x1, dim)
min!(dst, f2, x1, x2, dim)
min!(dst, f3, x1, x2, x3, dim)

minfdiff(f2, x, y)     # compute minimum of f2(x - y)
minfdiff(f2, x, y, dim)
minfdiff!(dst, f2, x, y, dim)

```

Note: when computing maximum/minimum along specific dimension, we use `max(x, (), dim)` and `min(x, (), dim)` instead of `max(x, dim)` and `min(x, dim)` to avoid ambiguities that would otherwise occur.

5.2 Generic folding

This package extends `foldl` and `foldr` for generic folding.

```
# suppose length(x) == 4
foldl(op, x)      # i.e. op(op(op(x[1], x[2]), x[3]), x[4])
foldr(op, x)      # i.e. op(x[1], op(x[2], op(x[3], x[4])))

foldl(Add(), x)   # sum over x from left to right
foldr(Add(), x)   # sum over x from right to left
```

You can also use functors to generate terms for folding.

```
foldl(op, f1, x)   # fold over f1(x) from left to right
foldr(op, f1, x)   # fold over f1(x) from right to left

foldl(op, f2, x1, x2) # fold over f2(x1, x2) from left to right
foldr(op, f2, x1, x2) # fold over f2(x1, x2) from right to left

foldl(op, f3, x1, x2, x3) # fold over f3(x1, x2, x3) from left to right
foldr(op, f3, x1, x2, x3) # fold over f3(x1, x2, x3) from right to left

foldl_fdiff(op, f, x, y) # fold over f(x - y) from left to right
foldr_fdiff(op, f, x, y) # fold over f(x - y) from right to left
```

You may also provide an initial value `s0` for folding.

```
foldl(op, s0, x)
foldr(op, s0, x)

foldl(op, s0, f1, x)
foldr(op, s0, f1, x)

foldl(op, s0, f2, x1, x2)
foldf(op, s0, f2, x1, x2)

foldl(op, s0, f3, x1, x2, x3)
foldr(op, s0, f3, x1, x2, x3)

foldl_fdiff(op, s0, f, x, y)
foldr_fdiff(op, s0, f, x, y)
```

The function `foldl` also supports reduction along a specific `dim`.

```
foldl(op, s0, x, dim)           # fold op over x along dimension dim
foldl(op, s0, f1, x, dim)      # fold op over f1(x) along dimension dim
foldl(op, s0, f2, x1, x2, dim) # fold op over f2(x1, x2) along dimension dim
foldl(op, s0, f3, x1, x2, x3, dim) # fold op over f3(x1, x2, x3) along dimension dim
foldl_fdiff(op, s0, f, x, y, dim) # fold op over f(x - y) along dimension dim

# the following statement write results to pre-allocated storage

foldl!(dst, op, s0, x, dim)
foldl!(dst, op, s0, f1, x, dim)
foldl!(dst, op, s0, f2, x1, x2, dim)
foldl!(dst, op, s0, f3, x1, x2, x3, dim)
foldl_fdiff!(dst, op, s0, f, x, y, dim)
```

5.3 Derived reduction functions

In addition to these basic reduction functions, we also define a set of derived reduction functions, as follows:

```
var(x)
var(x, dim)
var!(dst, x, dim)

std(x)
std(x, dim)
std!(dst, x, dim)

sumabs(x) # == sum(abs(x))
sumabs(x, dim)
sumabs!(dst, x, dim)

meanabs(x) # == mean(abs(x))
meanabs(x, dim)
meanabs!(dst, x, dim)

maxabs(x) # == max(abs(x))
maxabs(x, dim)
maxabs!(dst, x, dim)

minabs(x) # == min(abs(x))
minabs(x, dim)
minabs!(dst, x, dim)

sumsq(x) # == sum(abs2(x))
sumsq(x, dim)
sumsq!(dst, x, dim)

meansq(x) # == mean(abs2(x))
meansq(x, dim)
meansq!(dst, x, dim)

dot(x, y) # == sum(x .* y)
dot(x, y, dim)
dot!(dst, x, y, dim)

sumabsdiff(x, y) # == sum(abs(x - y))
sumabsdiff(x, y, dim)
sumabsdiff!(dst, x, y, dim)

meanabsdiff(x, y) # == mean(abs(x - y))
meanabsdiff(x, y, dim)
meanabsdiff!(dst, x, y, dim)

maxabsdiff(x, y) # == max(abs(x - y))
maxabsdiff(x, y, dim)
maxabsdiff!(dst, x, y, dim)

minabsdiff(x, y) # == min(abs(x - y))
minabsdiff(x, y, dim)
minabsdiff!(dst, x, y, dim)

sumsqdiff(x, y) # == sum(abs2(x - y))
sumsqdiff(x, y, dim)
```

```
sumsqdiff!(dst, x, y, dim)

meansqdiff(x, y) # == mean(abs2(x - y))
meansqdiff(x, y, dim)
meansqdiff!(dst, x, y, dim)
```

Although this is quite a large set of functions, the actual code is quite concise, as most of such functions are generated through macros (see `src/reduce.jl`)

In addition to the common reduction functions, this package also provides a set of statistics functions that are particularly useful in probabilistic or information theoretical computation, as follows

```
sumxlogx(x) # == sum(xlogx(x)) with xlogx(x) = x > 0 ? x * log(x) : 0
sumxlogx(x, dim)
sumxlogx!(dst, x, dim)

sumxlogy(x, y) # == sum(xlog(x,y)) with xlogy(x,y) = x > 0 ? x * log(y) : 0
sumxlogy(x, y, dim)
sumxlogy!(dst, x, y, dim)

entropy(x) # == - sumxlogx(x)
entropy(x, dim)
entropy!(dst, x, dim)

logsumexp(x) # == log(sum(exp(x)))
logsumexp(x, dim)
logsumexp!(dst, x, dim)

softmax!(dst, x) # dst[i] = exp(x[i]) / sum(exp(x))
softmax(x)
softmax!(dst, x, dim)
softmax(x, dim)
```

For `logsumexp` and `softmax`, special care is taken to ensure numerical stability for large `x` values, that is, their values will be properly shifted during computation (e.g. you can perfectly do `logsumexp([1000., 2000., 3000.])` with this package, while `log(sum(exp([1000., 2000., 3000.]))` would lead to overflow.)

5.4 Weighted Sum

Computation of weighted sum as below is common in practice.

$$\sum_{i=1}^n w_i x_i$$

$$\sum_{i=1}^n w_i f(x_i, \dots)$$

$$\sum_{i=1}^n w_i f(x_i - y_i)$$

NumericExtensions.jl directly supports such computation via `wsum` and `wsumfdiff`:

```
wsum(w, x) # weighted sum of x with weights w
wsum(w, f1, x1) # weighted sum of f1(x1) with weights w
wsum(w, f2, x1, x2) # weighted sum of f2(x1, x2) with weights w
wsum(w, f3, x1, x2, x3) # weighted sum of f3(x1, x2, x3) with weights w
wsumfdiff(w, f2, x, y) # weighted sum of f2(x - y) with weights w
```

These functions also support computing the weighted sums along a specific dimension:

```
wsum(w, x, dim)
wsum!(dst, w, x, dim)

wsum(w, f1, x1, dim)
wsum!(dst, w, f1, x1, dim)

wsum(w, f2, x1, x2, dim)
wsum!(dst, w, f2, x1, x2, dim)

wsum(w, f3, x1, x2, x3, dim)
wsum!(dst, w, f3, x1, x2, x3, dim)

wsumfdiff(w, f2, x, y, dim)
wsumfdiff!(dst, w, f2, x, y, dim)
```

Furthermore, `wsumabs`, `wsumabsdiff`, `wsumsq`, `wsumsqdiff` are provided to compute weighted sum of absolute values / squares to simplify common use:

```
wsumabs(w, x)                # weighted sum of abs(x)
wsumabs(w, x, dim)
wsumabs!(dst, w, x, dim)

wsumabsdiff(w, x, y)        # weighted sum of abs(x - y)
wsumabsdiff(w, x, y, dim)
wsumabsdiff!(dst, w, x, y, dim)

wsumsq(w, x)                # weighted sum of abs2(x)
wsumsq(w, x, dim)
wsumsq!(dst, w, x, dim)

wsumsqdiff(w, x, y)        # weighted sum of abs2(x - y)
wsumsqdiff(w, x, y, dim)
wsumsqdiff!(dst, w, x, y, dim)
```

5.5 Performance

The reduction and map-reduction functions are carefully optimized. In particular, several tricks lead to performance improvement:

- computation is performed in a cache-friendly manner;
- computation completes in a single pass without creating intermediate arrays;
- kernels are inlined via the use of typed functors;
- inner loops use linear indexing (with pre-computed offset);
- opportunities of using BLAS are exploited.

Generally, many of the reduction functions in this package can achieve $3x - 12x$ speed up as compared to the typical Julia expression.

We observe further speed up for certain functions:

- full reduction with `sumabs`, `sumsq`, and `dot` utilize BLAS level 1 routines, and they achieve $10x$ to $30x$ speed up.

- For `var` and `std`, we devise dedicated procedures, where computational steps are very carefully scheduled such that most computation is conducted in a single pass. This results in about $25x$ speedup.

Vector Norms and Normalization

This package provides functions for evaluating vector norms and normalizing vectors.

6.1 Vector Norm Evaluation

Synopsis

```

vnorm(x, p)           # compute L-p norm of vec(x)
vnorm(x, p, dim)      # compute L-p norm of x along dimension dim
vnorm!(r, x, p, dim)  # compute L-p norm along specific dimension and
                       # write results to r

vnormdiff(x, y, p)    # compute L-p norm of vec(x - y)
vnormdiff(x, y, p, dim) # compute L-p norm of x - y along dim
vnormdiff!(r, x, y, p, dim) # compute L-p norm of x - y along dim and write to r

```

Notes:

- For `vnormdiff` and `vnormdiff!`, `x` or `y` can be either an array or a scalar.
- When `p` is 1, 2, or `Inf`, specialized fast routines are used.

Examples

```

vnorm(x, 2)           # compute L-2 norm of x
vnorm(x, 2, 1)        # compute L-2 norm of each column of x
vnorm(x, Inf, 2)      # compute L-inf norm of each row of x
vnorm!(r, x, 2, 1)    # compute L-2 norm of each column, and write results to r

vnormdiff(x, 2.5, 2)  # compute L-2 norm of x - 2.5
vnormdiff(x, y, 1, 2) # compute L-1 norm of x - y for each column

```

6.2 Normalization

Normalizing a vector w.r.t L-p norm means to scale a vector such that the L-p norm of the vector becomes 1.

Synopsis

```

normalize(x, p)        # returns a normalized vector w.r.t. L-p norm
normalize!(x, p)       # normalize x w.r.t. L-p norm inplace
normalize!(r, x, p)    # write the normalized vector to a pre-allocated array r

```

```
normalize(x, p, dim)      # returns an array comprised of normalized vectors along dim
normalize!(x, p, dim)     # normalize vectors of x along dim inplace w.r.t. L-p norm
normalize!(r, x, p, dim) # write the normalized vectors along dim to r
```

Array Scanning

Scanning refers to the operation that applies a binary function recursively as follows:

Let x be an input vector and f be a binary operator, then the output y is a vector of the same length given by

$$y[1] = x[1]$$

$$y[i] = f(y[i - 1], x[i]), \forall i = 2, 3, \dots$$

For example, the function `cumsum` is a special case of array scanning.

This package provides both generic scanning functions, such as `scan`, `scan!`, `mapscan`, and `mapscan!`, as well as specialized scanning functions, such as `cumsum` and `cumsum!`, etc.

7.1 Generic Scanning

Synopsis

Let `op` be a binary operator, `f1`, `f2`, and `f3` be respectively unary, binary, and ternary functors.

```
scan(op, x)                # scan x using operator op
mapscan(op, f1, x)         # scan f1(x)
mapscan(op, f2, x, y)     # scan f2(x, y)
mapscan(op, f3, x, y, z)  # scan f3(x, y, z)

scan!(r, op, x)           # write the results of scanning of x to r
mapscan!(r, op, f1, x)
mapscan!(r, op, f2, x, y)
mapscan!(r, op, f3, x, y, z)

scan!(op, x)              # inplace scanning of x using operator op
```

Both `scan` and `scan!` accepts a dimension argument for scanning vectors along a specific dimension.

```
scan(op, x, dim)          # scan vectors of x along dim using operator op
mapscan(op, f1, x, dim)   # scan vectors of f1(x) along dim
mapscan(op, f2, x, y, dim) # scan vectors of f2(x, y) along dim
mapscan(op, f3, x, y, z, dim) # scan vectors of f3(x, y, z) along dim

scan!(r, op, x, dim)     # write the results of scanning of x to r
mapscan!(r, op, f1, x, dim)
mapscan!(r, op, f2, x, y, dim)
mapscan!(r, op, f3, x, y, z, dim)
```

```
scan!(op, x, dim)           # inplace scanning of x along dim
```

Note: `mapscan` and `mapscan!` use efficient ways to scan the terms without creating a temporary array such as `f1(x)`. Hence, `mapscan(op, f1, x)` is generally faster than `scan(op, f(x))`.

Examples

```
scan(Add(), x)             # equivalent to cumsum(x)
scan(Add(), x, 1)         # equivalent to cumsum(x, 1)
scan!(Add(), x)           # writing cumsum(x) inplace (i.e. back to x)
```

7.2 Common Scanning Functions

For some common scanning operations, we provide specific functions to simplify the use.

```
# cumsum uses Add() as scanning operator

cumsum(x)
cumsum!(r, x)
cumsum!(x)

cumsum(f1, x)
cumsum(f2, x, y)
cumsum(f3, x, y, z)

cumsum!(r, f1, x)
cumsum!(r, f2, x, y)
cumsum!(r, f3, x, y, z)

# cummax uses MaxFun() as scanning operator

cummax(x)
cummax!(r, x)
cummax!(x)

cummax(f1, x)
cummax(f2, x, y)
cummax(f3, x, y, z)

cummax!(r, f1, x)
cummax!(r, f2, x, y)
cummax!(r, f3, x, y, z)

# cummin uses MinFun() as scanning operator

cummin(x)
cummin!(r, x)
cummin!(x)

cummin(f1, x)
cummin(f2, x, y)
cummin(f3, x, y, z)

cummin!(r, f1, x)
cummin!(r, f2, x, y)
cummin!(r, f3, x, y, z)
```