percentile

syntax: `percentile(data, fraction)`

purpose: Computes a value that is equal to or larger than the indicated fraction of the values in the data. If a list of fractions is specified, the percentile is computed separately for each element in the list.

examples: ≫ `x = [1 2 3 4 5 6 7];`
≫ `percentile(x, .5)` ⇒ *ans: 4*
Sometimes the returned value will not be one of the value in the data set.
≫ `x = [1 2 3 4 5 6];`
≫ `percentile(x, .5)` ⇒ *ans: 3.5*
In this case, the answer 3.5 is somewhat arbitrary; any value in the range 3 to 4 could have been chosen. In such situations PERCENTILE computes the mean of the two values that bracket the appropriate range.

Note that if you have a small number of samples, the percentile may not be very precise. To illustrate, we'll generate 10 samples from a standard normal distribution.
≫ `x = normal(10, 0, 1);`
The 5% and 95% percentiles are at the extremes of the data:
≫ `percentile(x, [.05, .95])`
ans: -1.4060 1.1900
Going even more extreme, to 2.5% and 97.5% gives the same result:
≫ `percentile(x, [.025, .975])`
ans: -1.4060 1.1900

When the data set is large, the extreme percentiles are much more meaningful.
≫ `x = normal(10000, 0, 1);`
≫ `percentile(x, [.05, .95])` ⇒ *ans: -1.66 1.67*
≫ `percentile(x, [.025, .975])` ⇒ *ans: -1.96 2.02*