

# FpFnPercentages

## FP%/FN% Percentages

The main system used to measure spam-filtering accuracy in [SpamAssassin](#) is the "FP%/FN% percentages" system.

It's quite simple. First, you scan a corpus of hand-classified mail (see [HandClassifiedCorpora](#)) to get 4 figures:

```
nspam = number of known-to-be-spam messages in the corpus
nham  = number of known-to-be-ham (nonspam) messages in the corpus
fp    = number of ham messages incorrectly marked as spam
fn    = number of spam messages incorrectly marked as ham
```

fp is so named because it's more commonly and concisely called a [FalsePositive](#), and fn a [FalseNegative](#).

Next, perform this calculation:

```
FP% = (fp / nham) * 100
FN% = (fn / nspam) * 100
```

and you have two numbers that simply, concisely, and comprehensively describe the accuracy and performance of the filter.

For example, let's say we do a test as follows:

```
nspam = 1000
nham  = 1500
fp    = 2
fn    = 30
```

the FP% and FN% work out as  $(2 / 1500) * 100 = 0.1333\%$  and  $(30 / 1000) * 100 = 3.0\%$  respectively.

The 'STATISTICS.txt' files distributed with [SpamAssassin](#) versions since about 2.30 include this data, measuring the ruleset's accuracy against a validation ruleset:

```
# SUMMARY for threshold 5.0:
# Correctly non-spam: 29443 99.97%
# Correctly spam:    27220 97.53%
# False positives:   9 0.03%
# False negatives:   688 2.47%
# TCR(l=50): 24.523726 SpamRecall: 97.535% SpamPrec: 99.967%
```

As you can see, FP% and FN% get pride of place in the measurement scheme.

See also [MeasuringAccuracy](#) for other methods.