PredictiveValue

Predictive Value

This is a scheme for MeasuringAccuracy, used in Renaissoft's Maia Mailguard among others. quoting a mail from Robert Le Blanc (subject 'Re: Brightmail', Tue 30 Nov 2004 10:06:06 -0800, sent to the SpamAssassin users list):

'Gray, Richard wrote:

```
| Can anyone shed any light on how Brightmail achieves the rather
| impressive statistics it is quoting, or do you think it is just smoke
| and mirrors?
```

The first thing to note about performance statistics with regard to a spam filter (no matter who makes it) is that there's no "standard" way to measure such a thing. Each manufacturer can use his own definition of "accuracy" or "efficiency", and tailor his formula to produce the most flattering result. Until you know what that formula is, however, you can't readily compare one product with another.

The closest thing to a "standard" way of measuring a spam filter's effectiveness is the scientific model that medical researchers use for diagnostic tests. Even so, there are five separate tests, not just one:

```
PPV = spam / (spam + FP)
NPV = ham / (ham + FN)
Sensitivity = spam / (spam + FN)
Specificity = ham / (ham + FP)
Efficiency = (spam + ham) / (spam + ham + FP + FN)
```

PPV is the Positive Predictive Value. If the filter says it's spam, how likely is it to actually be spam?

NPV is the Negative Predictive Value. If the filter says it's ham, how likely is it to actually be ham?

Sensitivity is the "true positive" rate. If it's actually spam, how likely is the filter to say it's spam?

Specificity is the "true negative" rate. If it's actually ham, how likely is the filter to say it's ham?

Efficiency is the ratio of true positives and true negatives to total mail items processed--that is, the percentage of mail that was correctly classified. This is what most people expect a vendor's claim to represent.

Needless to say, these five tests will give you five different statistics. On my SpamAssassin setup, for instance, my current stats look like this:

```
PPV = 99.68%
NPV = 98.07%
Sensitivity = 99.64%
Specificity = 98.23%
Efficiency = 99.43%
```

Now, if I were being honest about how well SpamAssassin has been working for me, I'd probably quote the Efficiency figure (99.43%), since I consider that the most comprehensive and realistic estimate of the filter's overall performance.

On the other hand, if I were selling this product and wanted to dazzle you with the most impressive statistic available, I'd cherry-pick the PPV figure (99.68%) and feed that number to the marketing department.'

As an editor's comment here, I (jm) would like to point out that in fact there *are* 'standard' ways to measure effectiveness; see MeasuringAccuracy for a few more. 😉