Probing: Extracting Class Probability from a Classifier

John Langford

(with Bianca Zadrozny)

Method

Observation 1: $c(x) = 1 \Rightarrow P(y = 1|x) > 0.5$

For importance-weighted classification:

If cost of false negative $= w$ and cost of false positive $= 1$ then $c(x) = 1 \Rightarrow P(y = 1|x) > \frac{1}{1+w}$

Basic guarantee: Good classification performance $\Rightarrow$small l_1 error in probability estimate.

Other Methods

1. Fit a sigmoid to a margin. (+fast -margin specific -no performance guarantee)

2. Count bagging proportion. (-result unrelated to $P(y|x)$ in general)

3. Classifier native. (+fast -classifier often terribly calibrated)

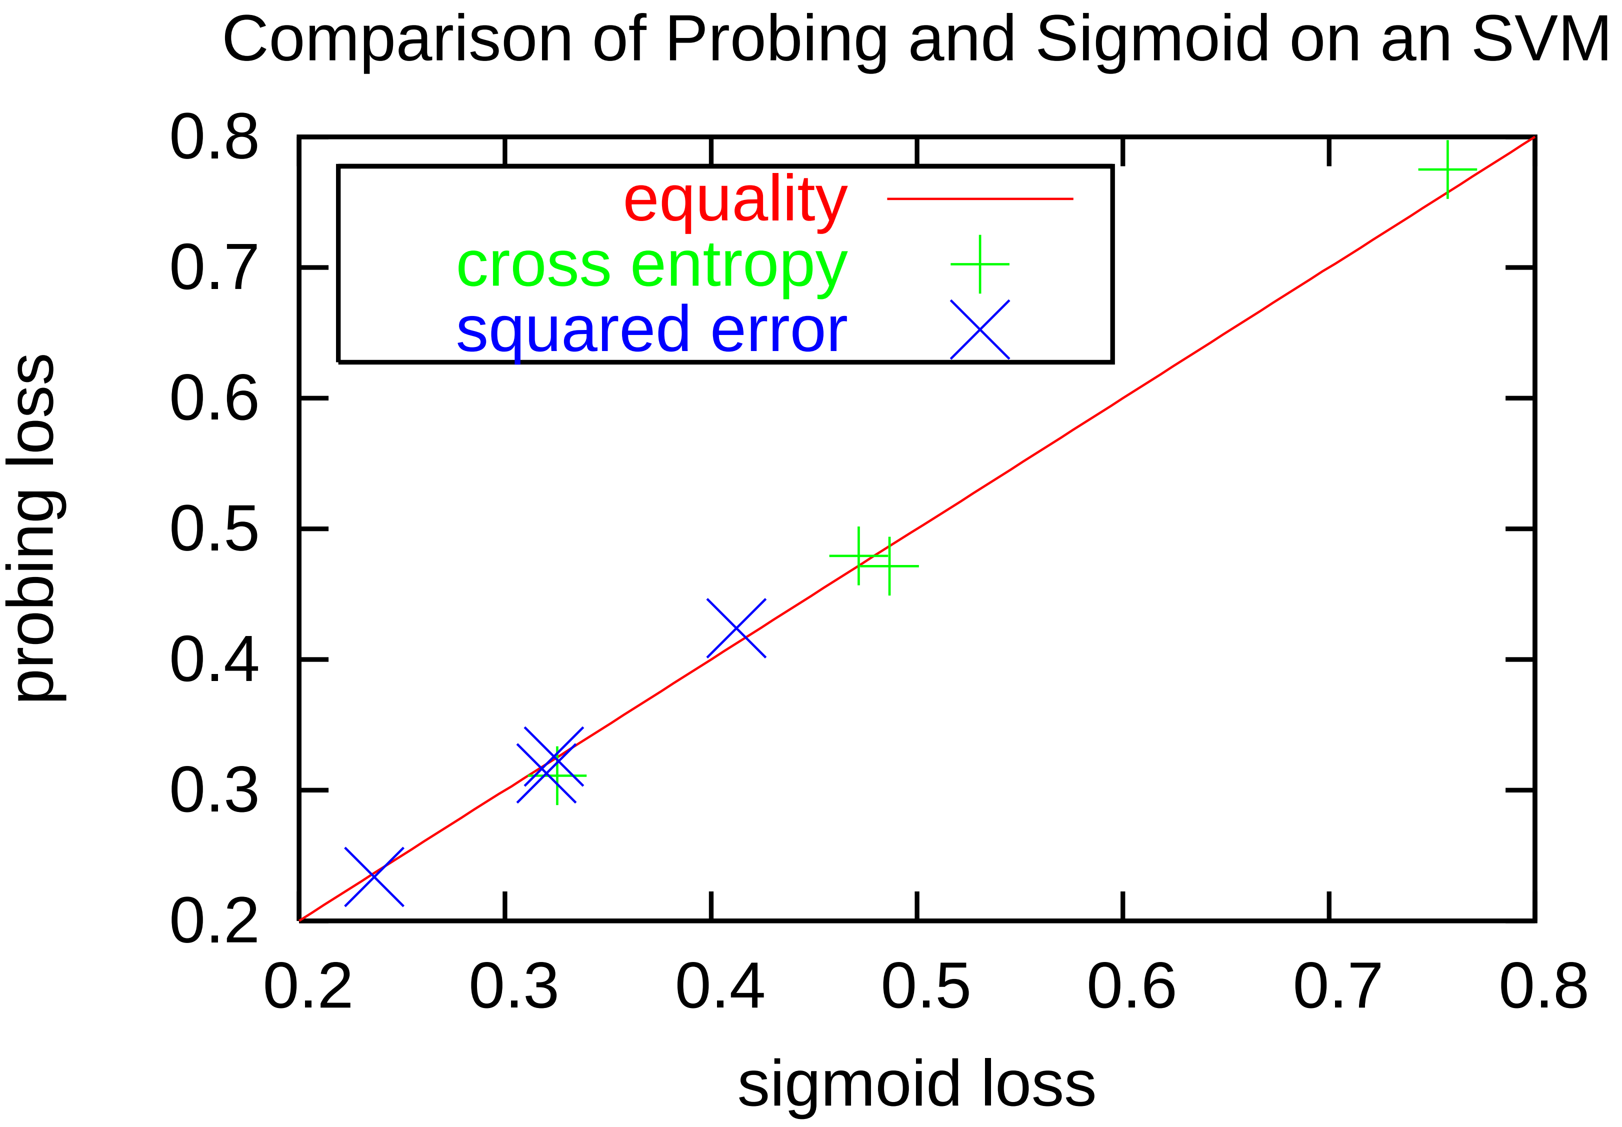
Comparison of Probing and Sigmoid on an SVM
equality
cross entropy
squared error
probing loss
sigmoid loss
0.2
0.3
0.4
0.5
0.6
0.7
0.8

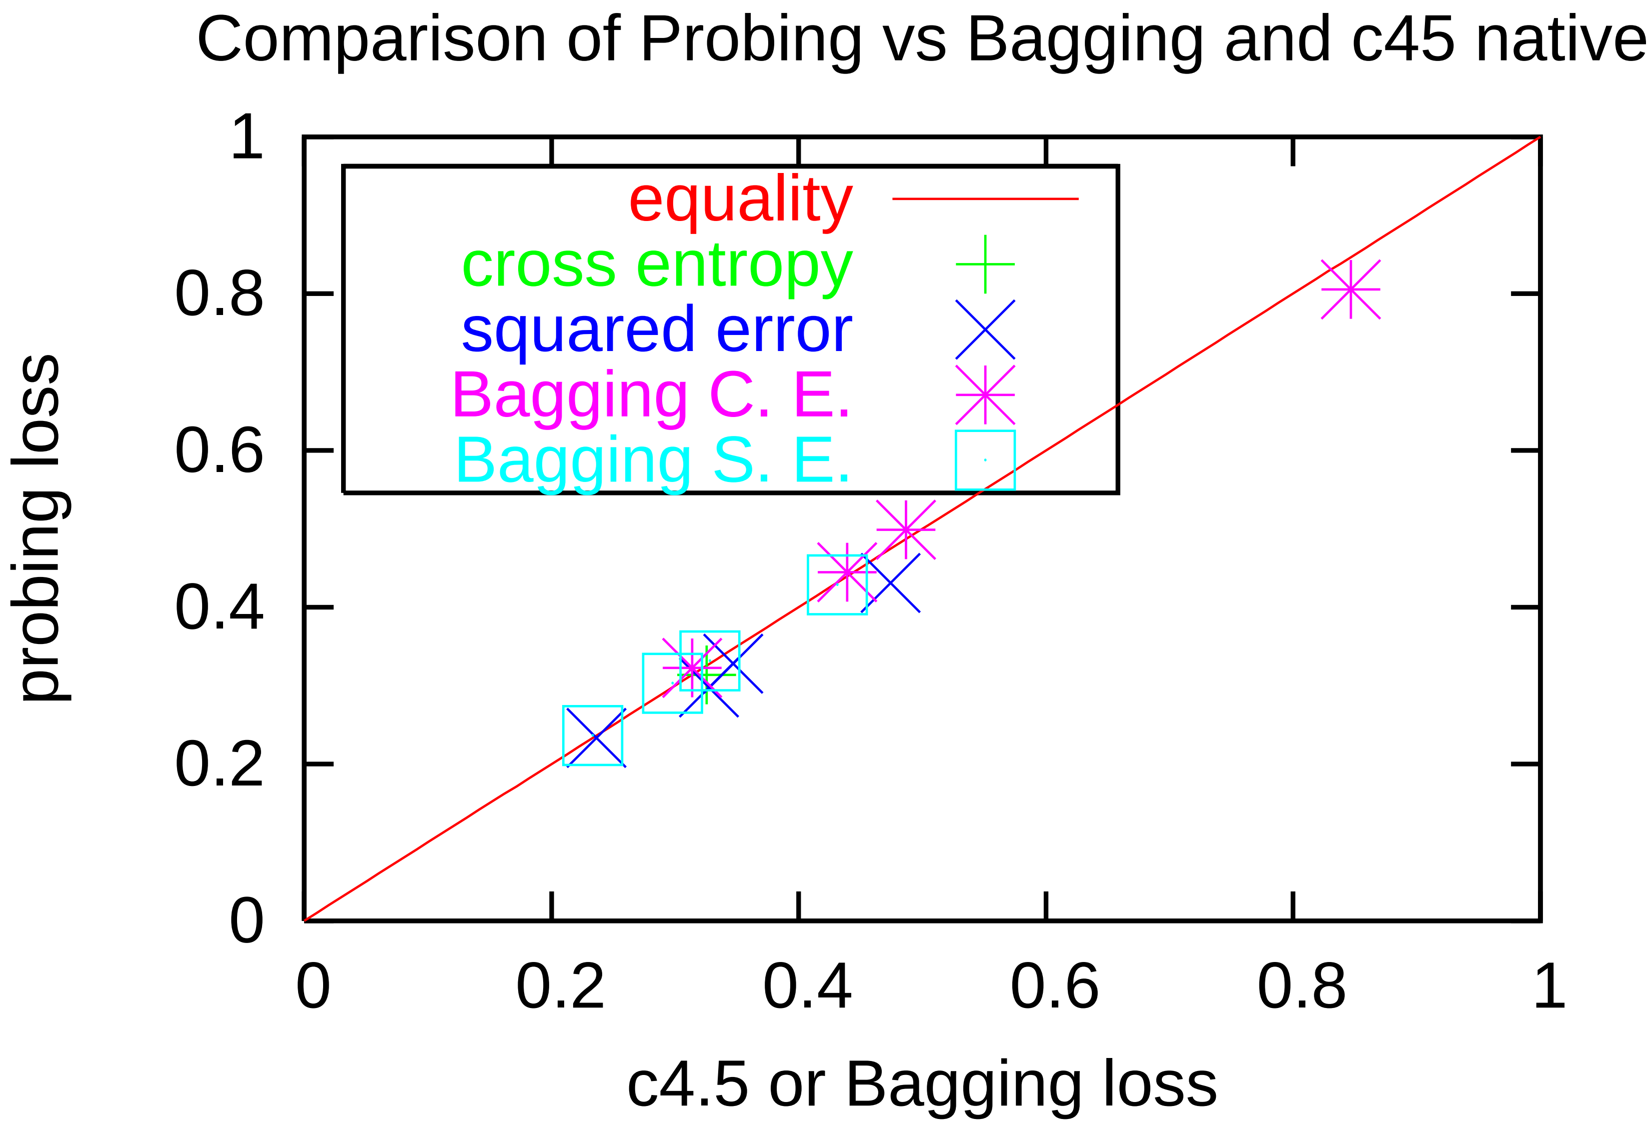
Comparison of Probing vs Bagging and c45 native
equality
cross entropy
squared error
Bagging C. E.
Bagging S. E.
probing loss
c4.5 or Bagging loss
0
0.2
0.4
0.6
0.8
1

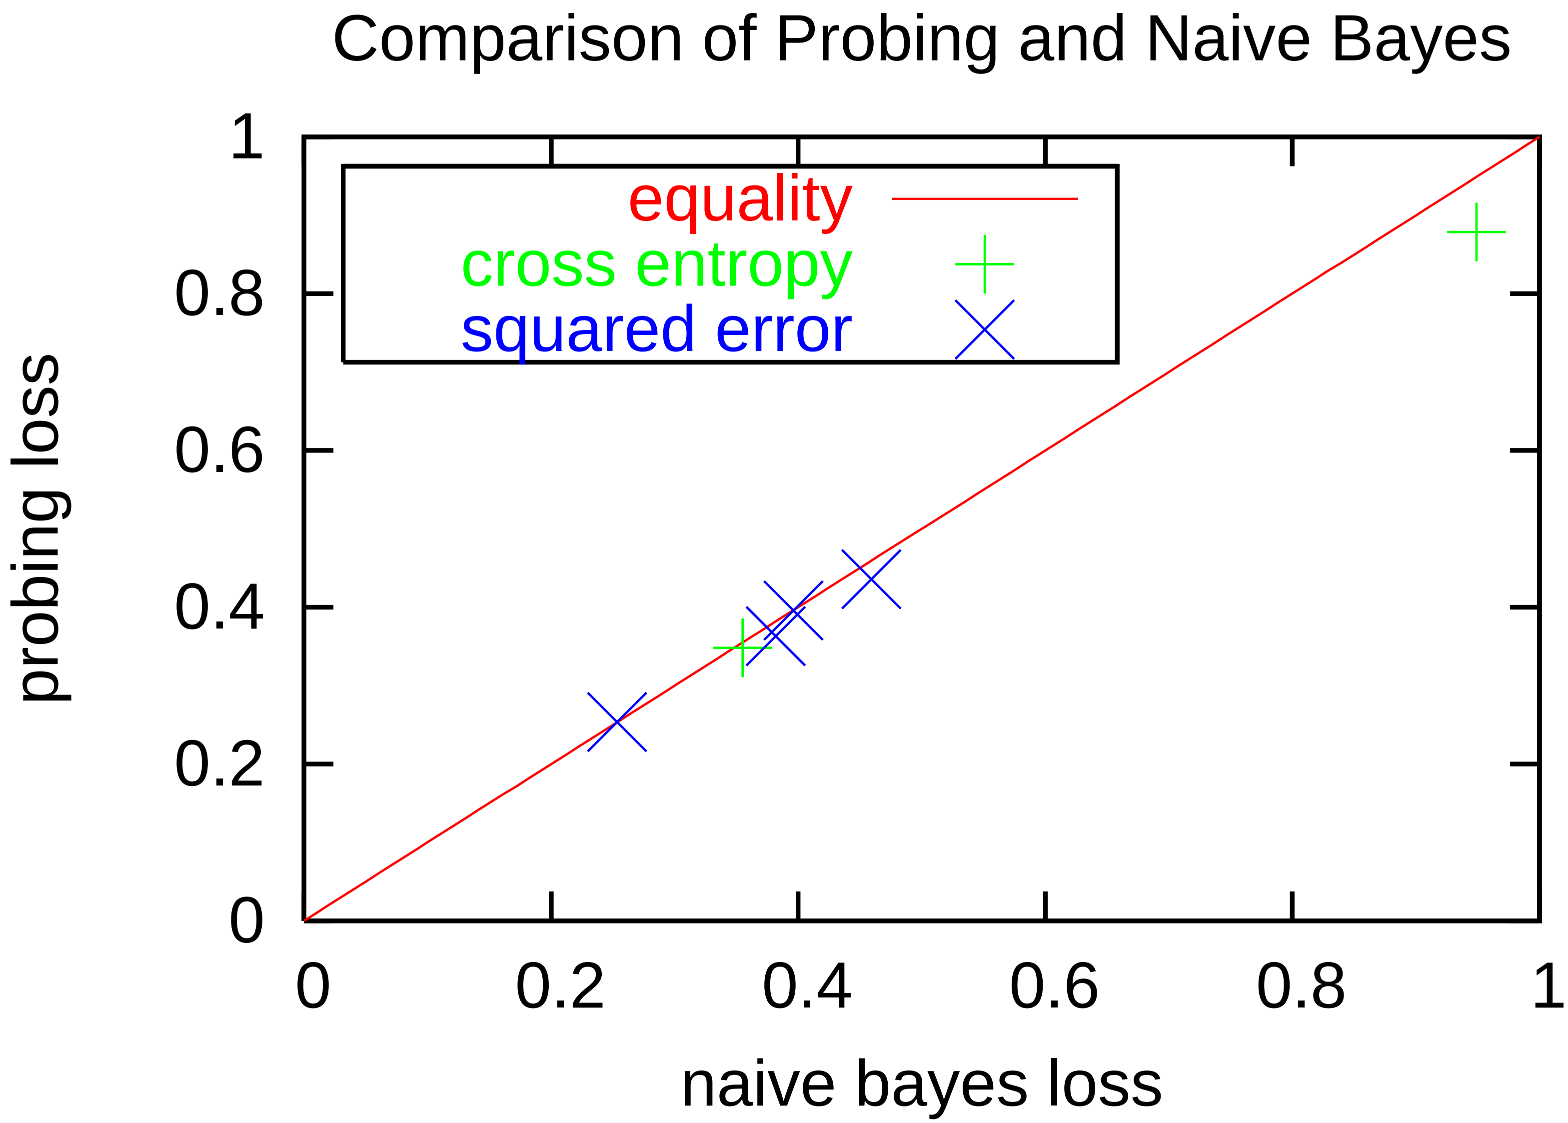
Comparison of Probing and Naive Bayes
equality
cross entropy
squared error
probing loss
naive bayes loss
0
0.2
0.4
0.6
0.8
1

Implications

1. Classification error rate $\epsilon \Rightarrow$ squared error loss $\leq \epsilon$.

2. No bound holds for cross entropy.

3. No direct bound on error rate for ranking.