Precision-Recall-Gain Curves

PR Analysis Done Right

Peter Flach, Meelis Kull @ bristol.ac.uk

1. Precision-Recall analysis and F-scores require proper treatment of their **harmonic scale** — arithmetic averages or linear expectations of F-scores etc are incoherent.

2. Precision-Recall-Gain curves properly linearise the quantities involved and their **area** is meaningful as an aggregate performance score.

3. These things matter **in practice** as AUPR can easily favour worse-performing models, unlike AUPRG.

4. Using PRG curves we can identify all F_β-optimal thresholds for any β in a single **calibration procedure**.

University of BRISTOL

INTELLIGENT SYSTEMS LAB

www.cs.bris.ac.uk/~flach/PRGcurves

chist-era EPSRC Engineering and Physical Sciences Research Council

REFRAME

ROC curve

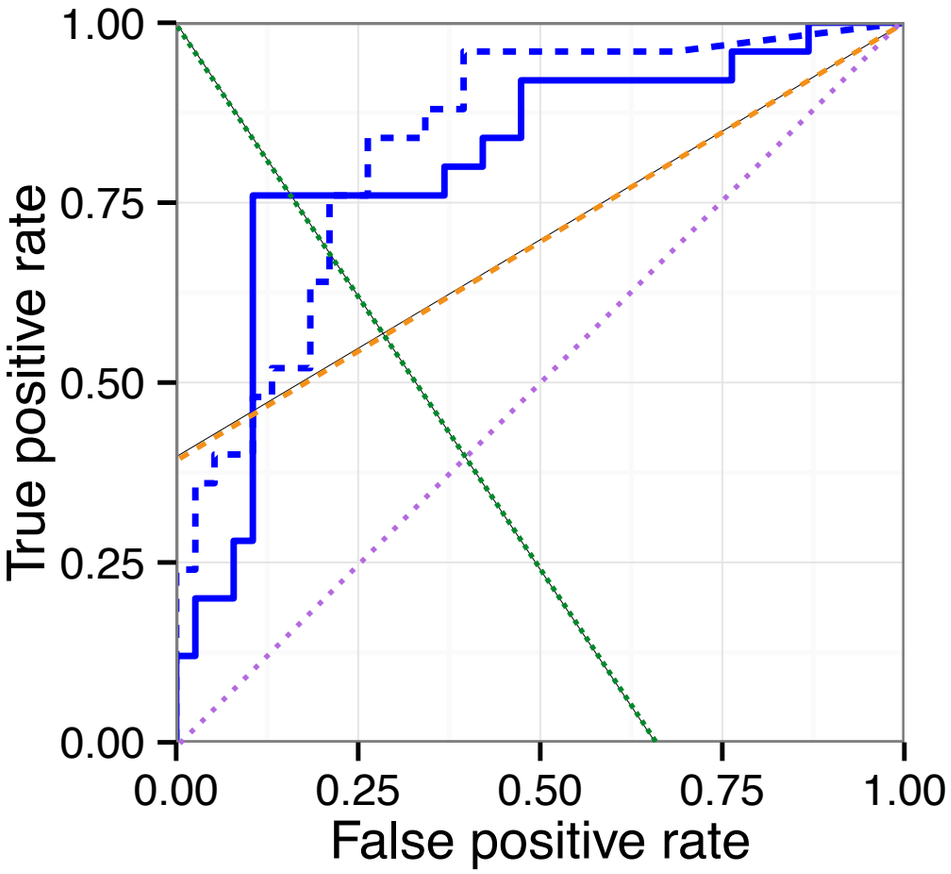

$$\mathbb{E}[acc] = \int_0^1 acc\, d\, rate = \pi(1-\pi)(2AUROC - 1) + 1/2$$

random

baseline

P=R

PR curve

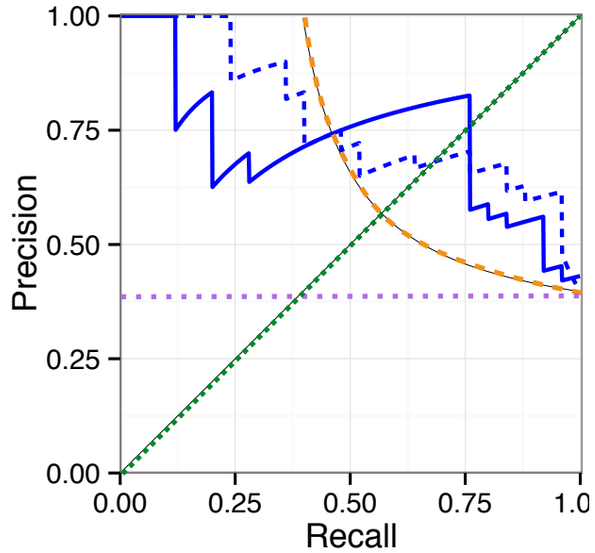

Precision-Recall-Gain curve

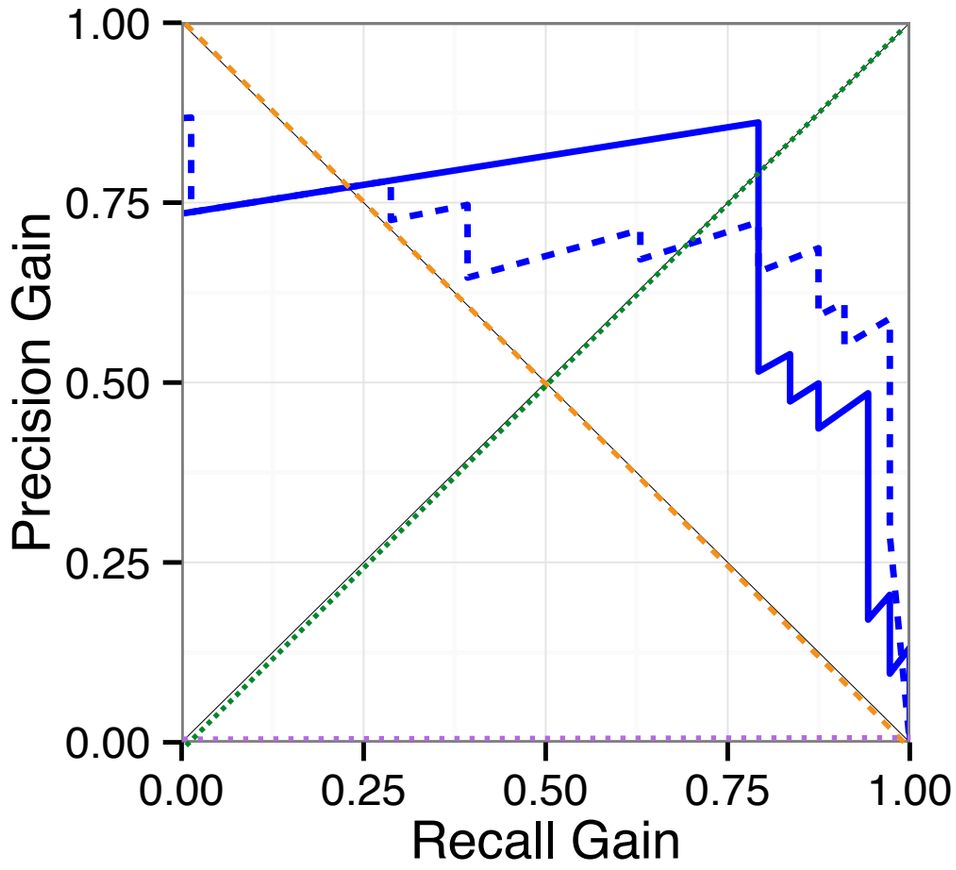

$$\mathbb{E}[FG_1] = \frac{AUPRG/2 + 1/4 - \pi(1-{y_0}^2)/4}{1 - \pi(1-y_0)}$$

University of BRISTOL

INTELLIGENT SYSTEMS LAB

Precision-Recall-Gain Curves

Peter Flach, Meelis Kull @ bristol.ac.uk

www.cs.bris.ac.uk/~flach/PRGcurves

chist-era

EPSRC Engineering and Physical Sciences Research Council

REFRAME

Accuracy calibration

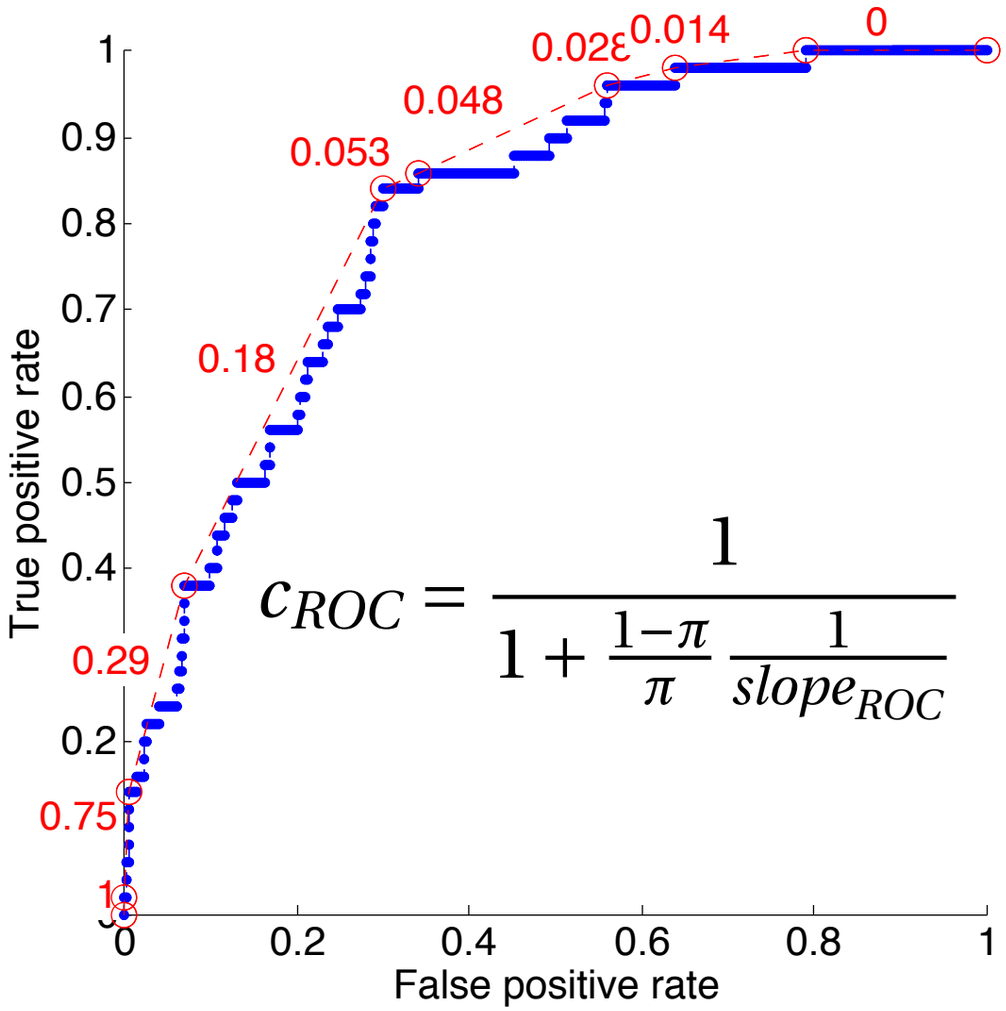

$$c_{ROC} = \frac{1}{1 + \frac{1-\pi}{\pi} \frac{1}{slope_{ROC}}}$$

F-score calibration

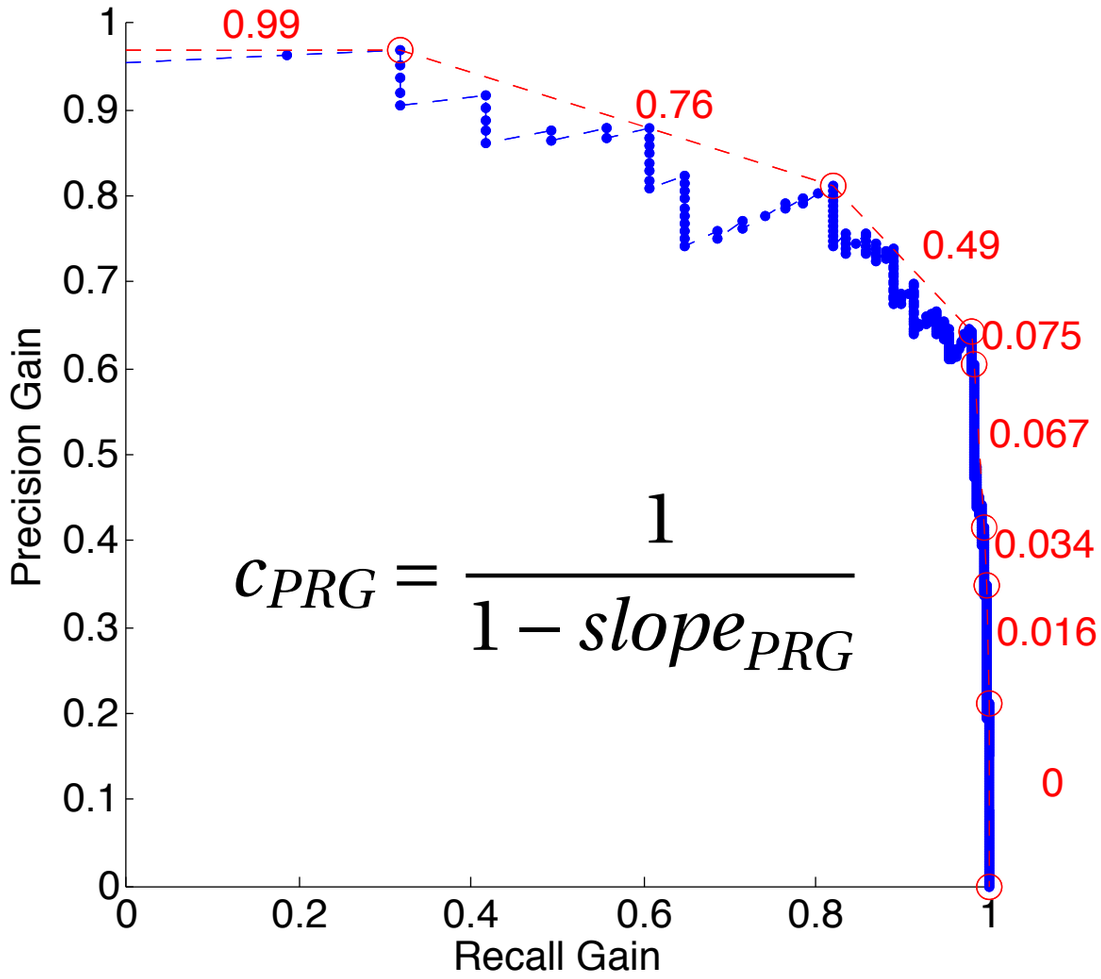

$$c_{PRG} = \frac{1}{1 - slope_{PRG}}$$

Precision-Recall-Gain Curves

Peter Flach, Meelis Kull @ bristol.ac.uk

www.cs.bris.ac.uk/~flach/PRGcurves

University of BRISTOL

INTELLIGENT SYSTEMS LAB

chist-era

EPSRC Engineering and Physical Sciences Research Council

REFRAME