
Falling Rule Lists (Supplement)

Fulton Wang
CSAIL and Sloan, MIT
fultonw@mit.edu

Cynthia Rudin
CSAIL and Sloan, MIT
rudin@mit.edu

Number of Rules Returned by Rule Miner

The nFoil based methods performed relatively poorly in part because the number of rules returned by nFoil was much smaller than the number of rules returned by FPGrowth, thus overly restricting the hypothesis space for such methods. We report in Table 1 the average, over the 5-fold cross validation we performed, of the number of rules mined by nFoil and FPGrowth, for every dataset considered in the paper.

Rule Miner	Readmissions	Spam	Mamm	Breast	Cars
nFoil	7	24	8	5	25
FPGrowth	103	550	43	147	190

Table 1: Number of rules mined for each dataset.

Running Times

We also report in Table 2 for each dataset the running times needed to fit a Falling Rule list on the entire dataset, as well as the size and number of features in it.

Dataset	n	p	running time (sec)
Readmission	7944	33	35
Spam	4601	57	102
Mamm	961	13	35
Breast	683	26	42
Cars	1728	21	39

Table 2: Running time for each dataset