

Overview Data Code Discussion Leaderboard Datasets Rules

Spearman's correlation score: 0.991141375982  
Kolmogorov-Smirnov test: KS-stat = 0.408810 p-value = 0.000e+00  
IMHO, it doesn't make sense to use KS-Test, if N is large.

Sergey Bryansky • (40th in this Competition) • 2 years ago • Options • Report • Reply

 2

In our ML world we have only three good techniques NN, GBM, and SVM;) I also tuned SVM with rbf kernel on original dataset it got .890+ 4-stratified-fold CV. [I used gpu implementation by Xtra Computing Group. https://github.com/Xtra-Computing/thundersvm](#)

Tilii Topic Author • (2931st in this Competition) • 2 years ago • Options • Report • Reply

 0

I've been meaning to give thundersvm a try, but have to adjust my existing scripts for grid searching. Good to know that it gets close, so I am fairly certain it will end up helping the ensemble.

antklen • 2 years ago • Options • Report • Reply

 2

That's interesting, how long does it take to train rbf ThunderSVM on that data?

Sergey Bryansky • (40th in this Competition) • 2 years ago • Options • Report • Reply

 1

It depends on kernel parameters and regularization. But on average it takes about 5 minutes per fold on 1080Ti card.

antklen • 2 years ago • Options • Report • Reply

 0

Not bad) thank you