

Learning curve extrapolation techniques across extrapolation settings

Lionel Kielhöfer¹, Felix Mohr², Jan N. van Rijn¹

¹ LIACS, Leiden University

² Universidad de La Sabana



Universiteit
Leiden
The Netherlands

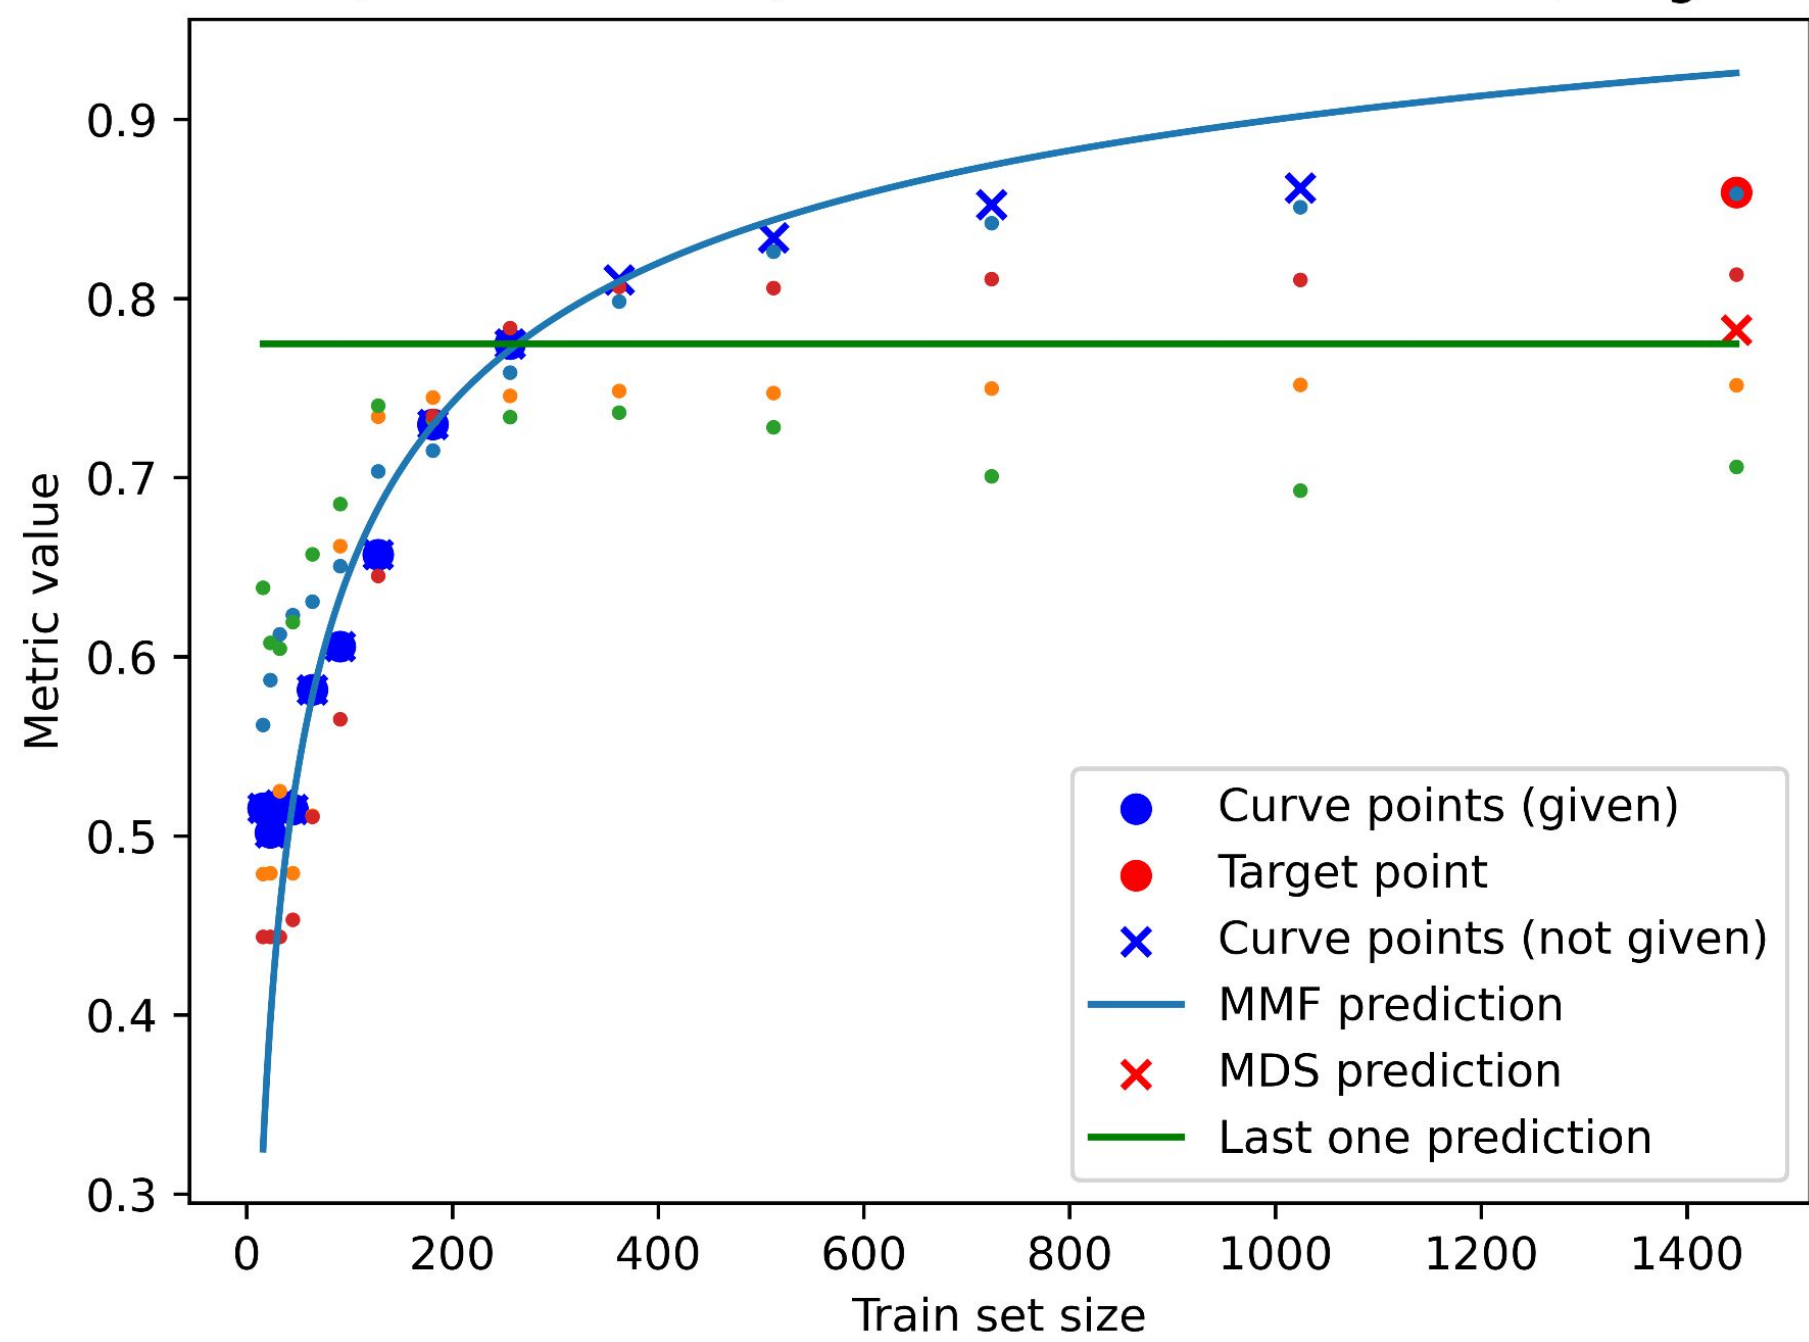


In a nutshell: Learning curve extrapolation techniques perform different depending on the problem setting (size of curve that is already explored, prediction target, etc)


Goal: Compare various learning curve extrapolation techniques across problem settings

Techniques: MDS (Leite and Brazdil, 2005, meta-learning based), MMF (Morgan, Mercer, Flodin, function-based), Best Last (baseline)

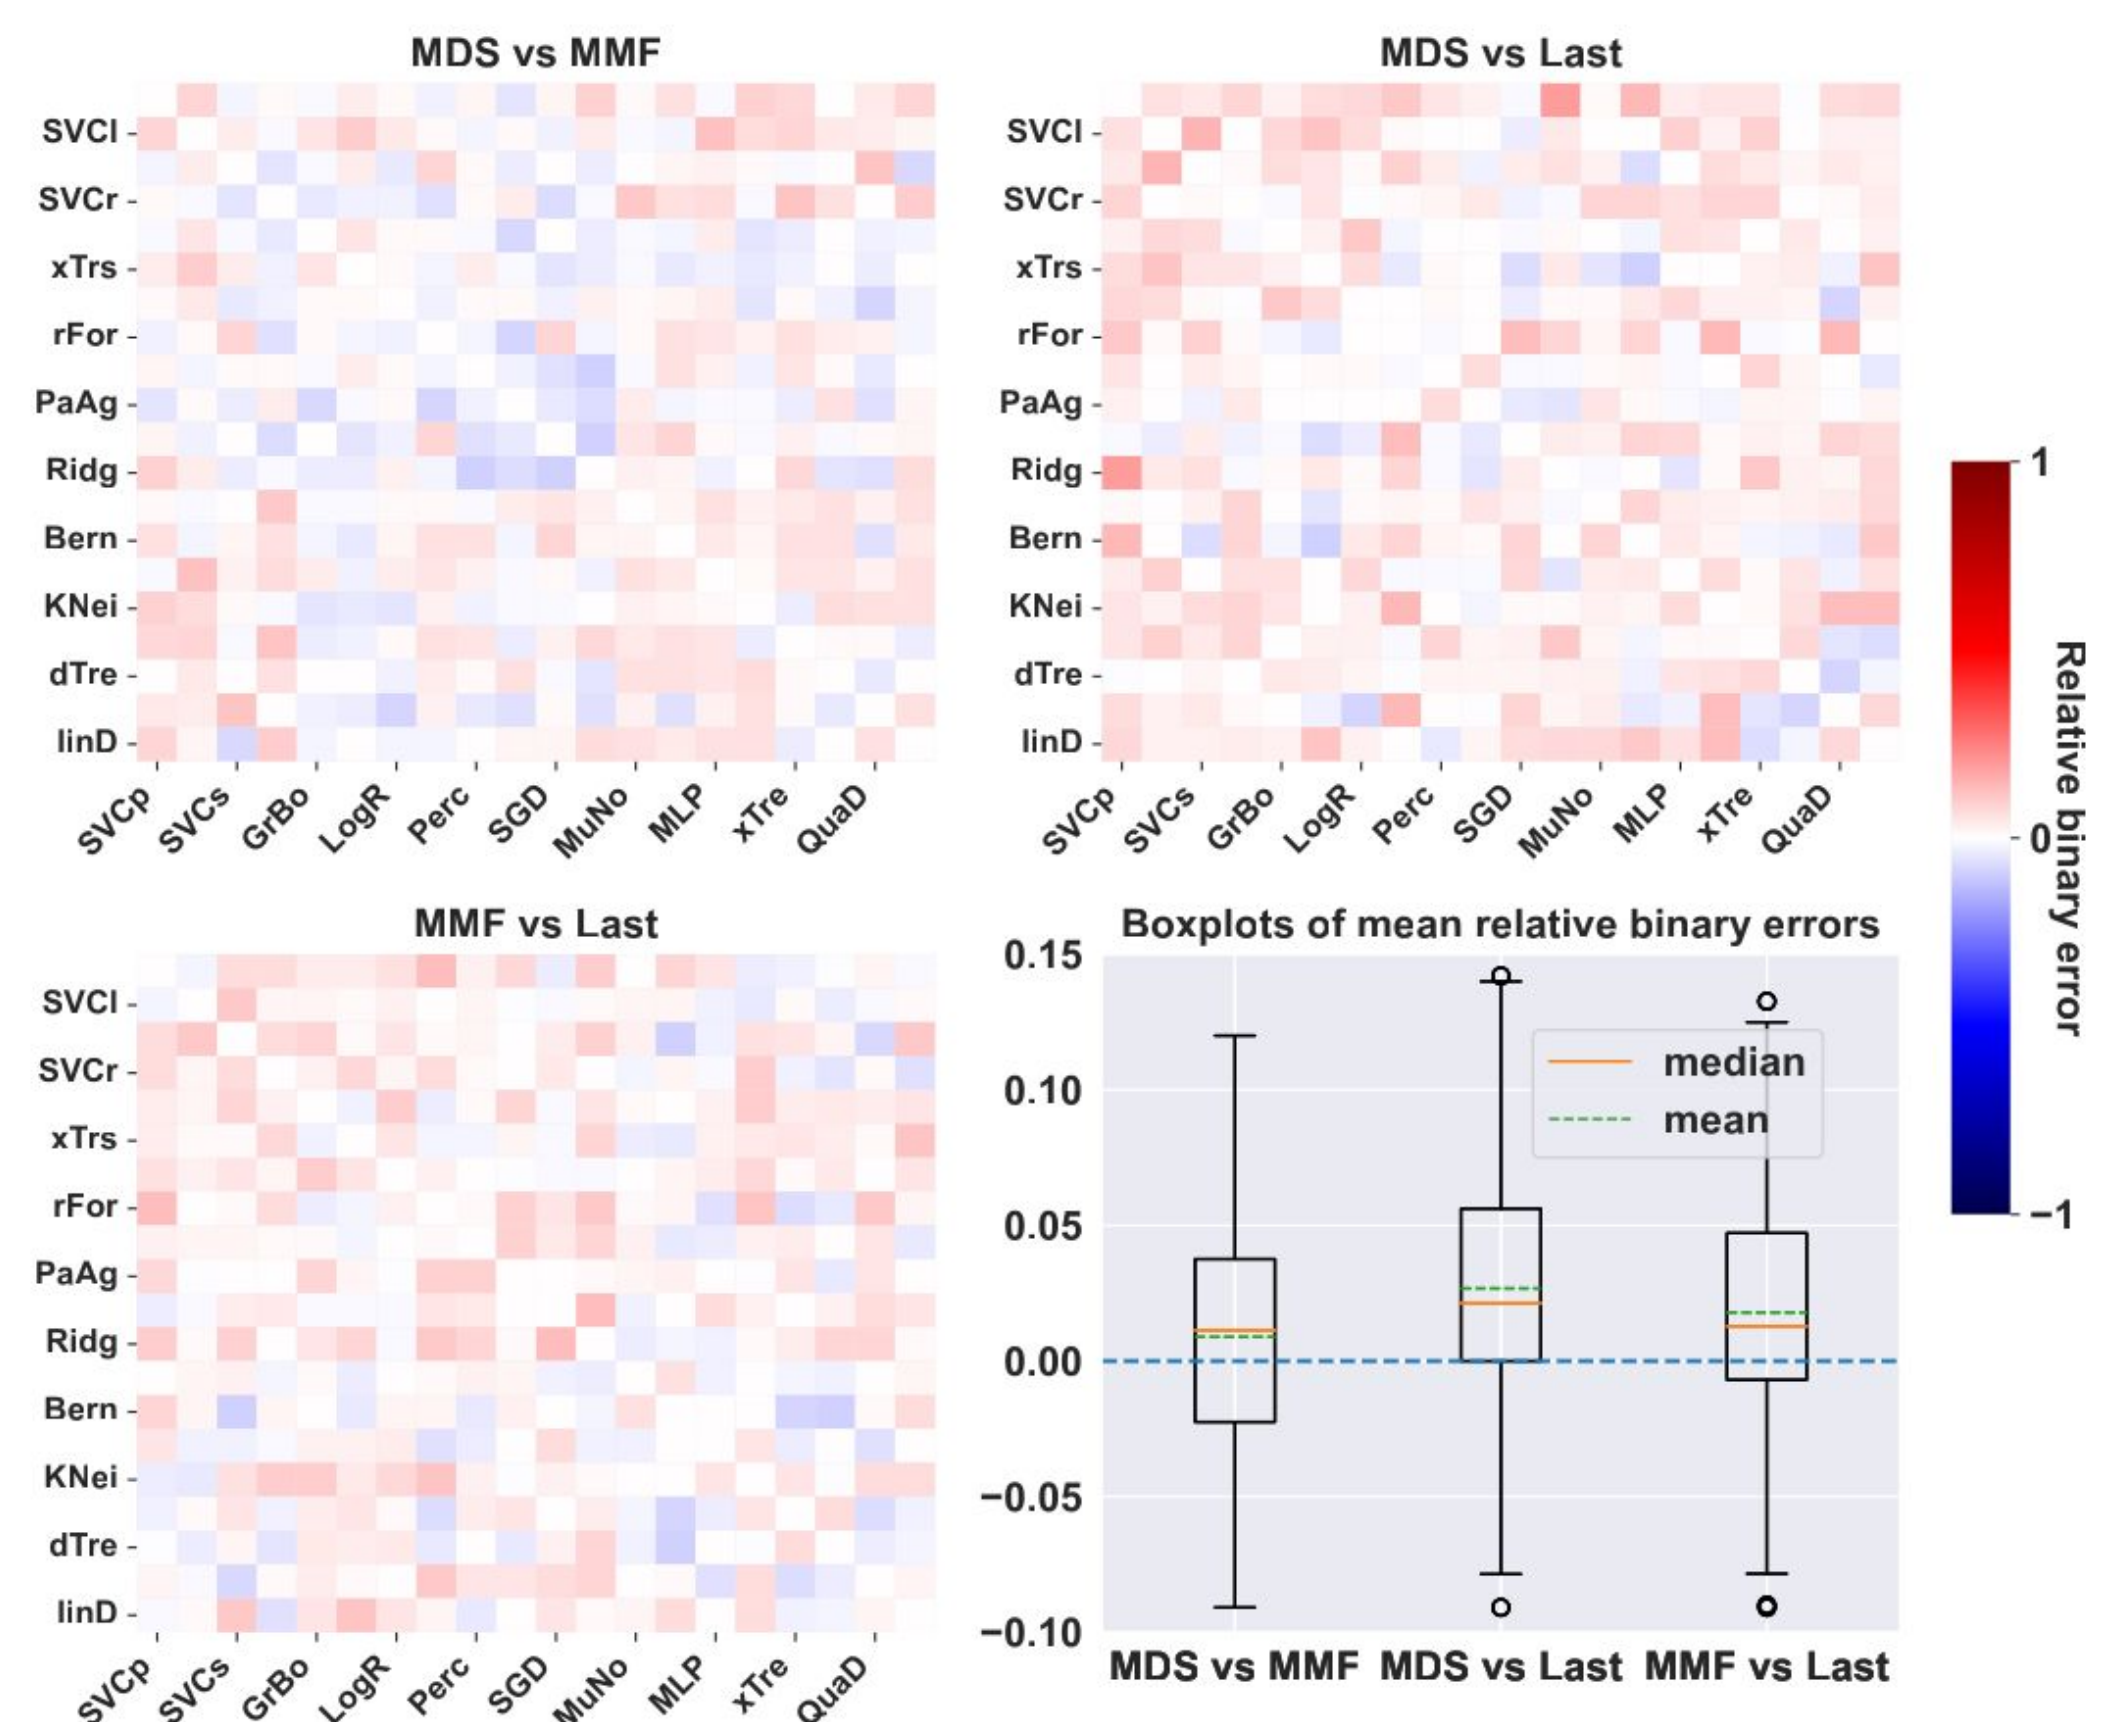
Dataset: 3, Learner: SVCs, Last anchor in window: 256, Target: 1448



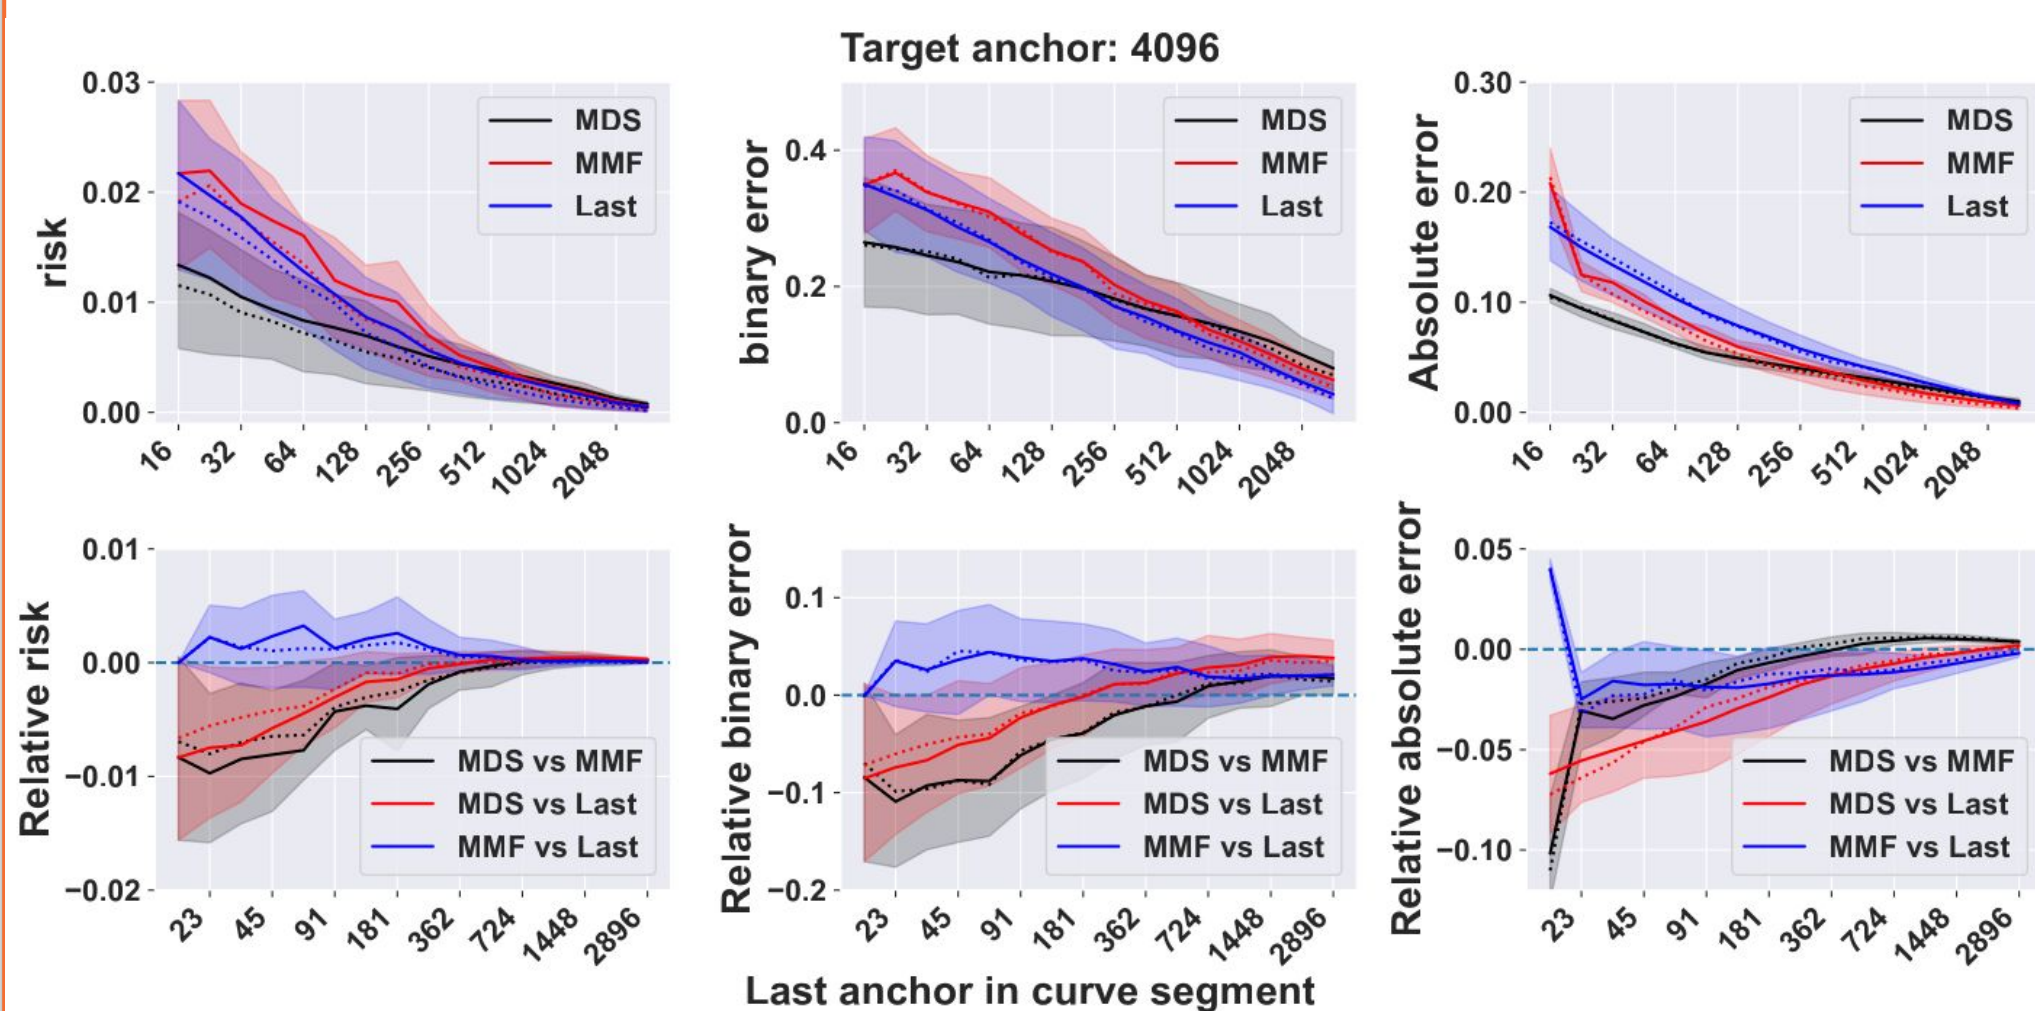
RQ1: How do the methods perform when we pick an arbitrary curve size and an arbitrary dataset size? (various classifiers)



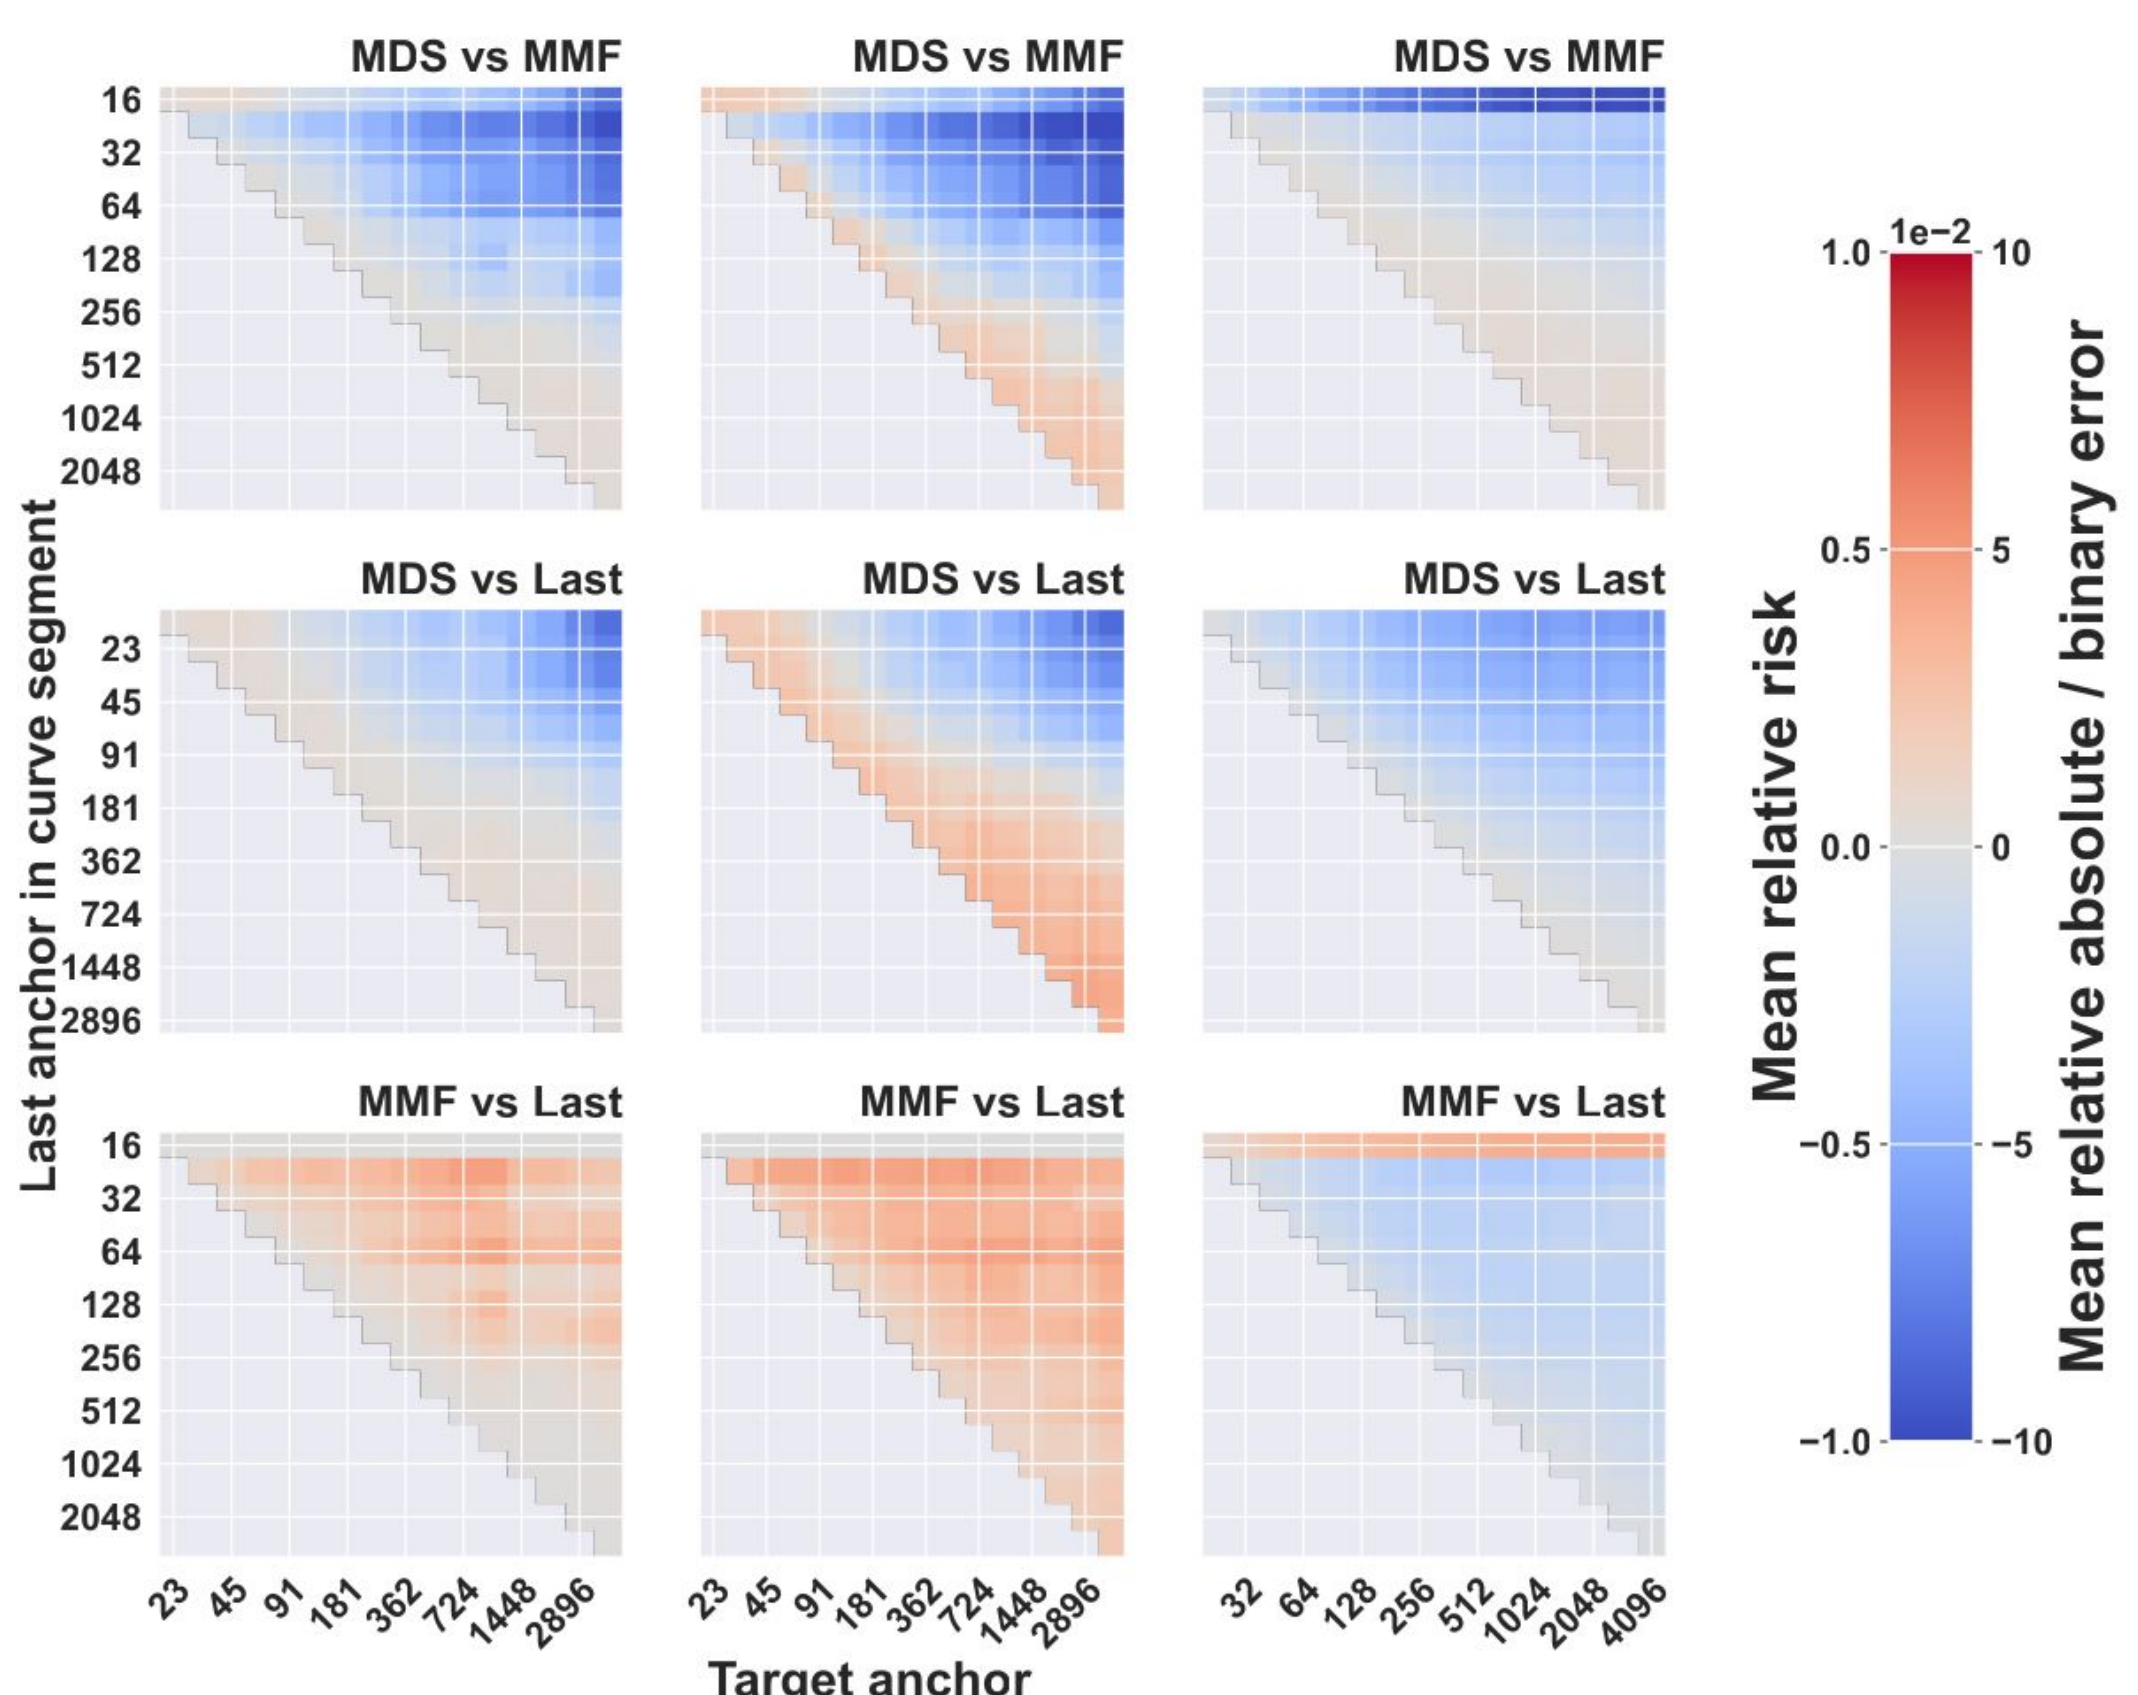
Last anchor in curve segment: 724, Target anchor: 4096



RQ2: How do the methods perform across various learning curve sizes?



RQ3: How do the methods perform across various learning curve sizes and target sizes?



Conclusions:

- When little data has been explored (small segment of learning curve is available), MDS is superior
- When the learning curve is extrapolated further, both MMF and the baseline perform better
- MMF does generally not outperform the baseline