


Near-linear scaling; 1M subjects in 1.4 s



Peak memory 153 MB at 1M subjects

