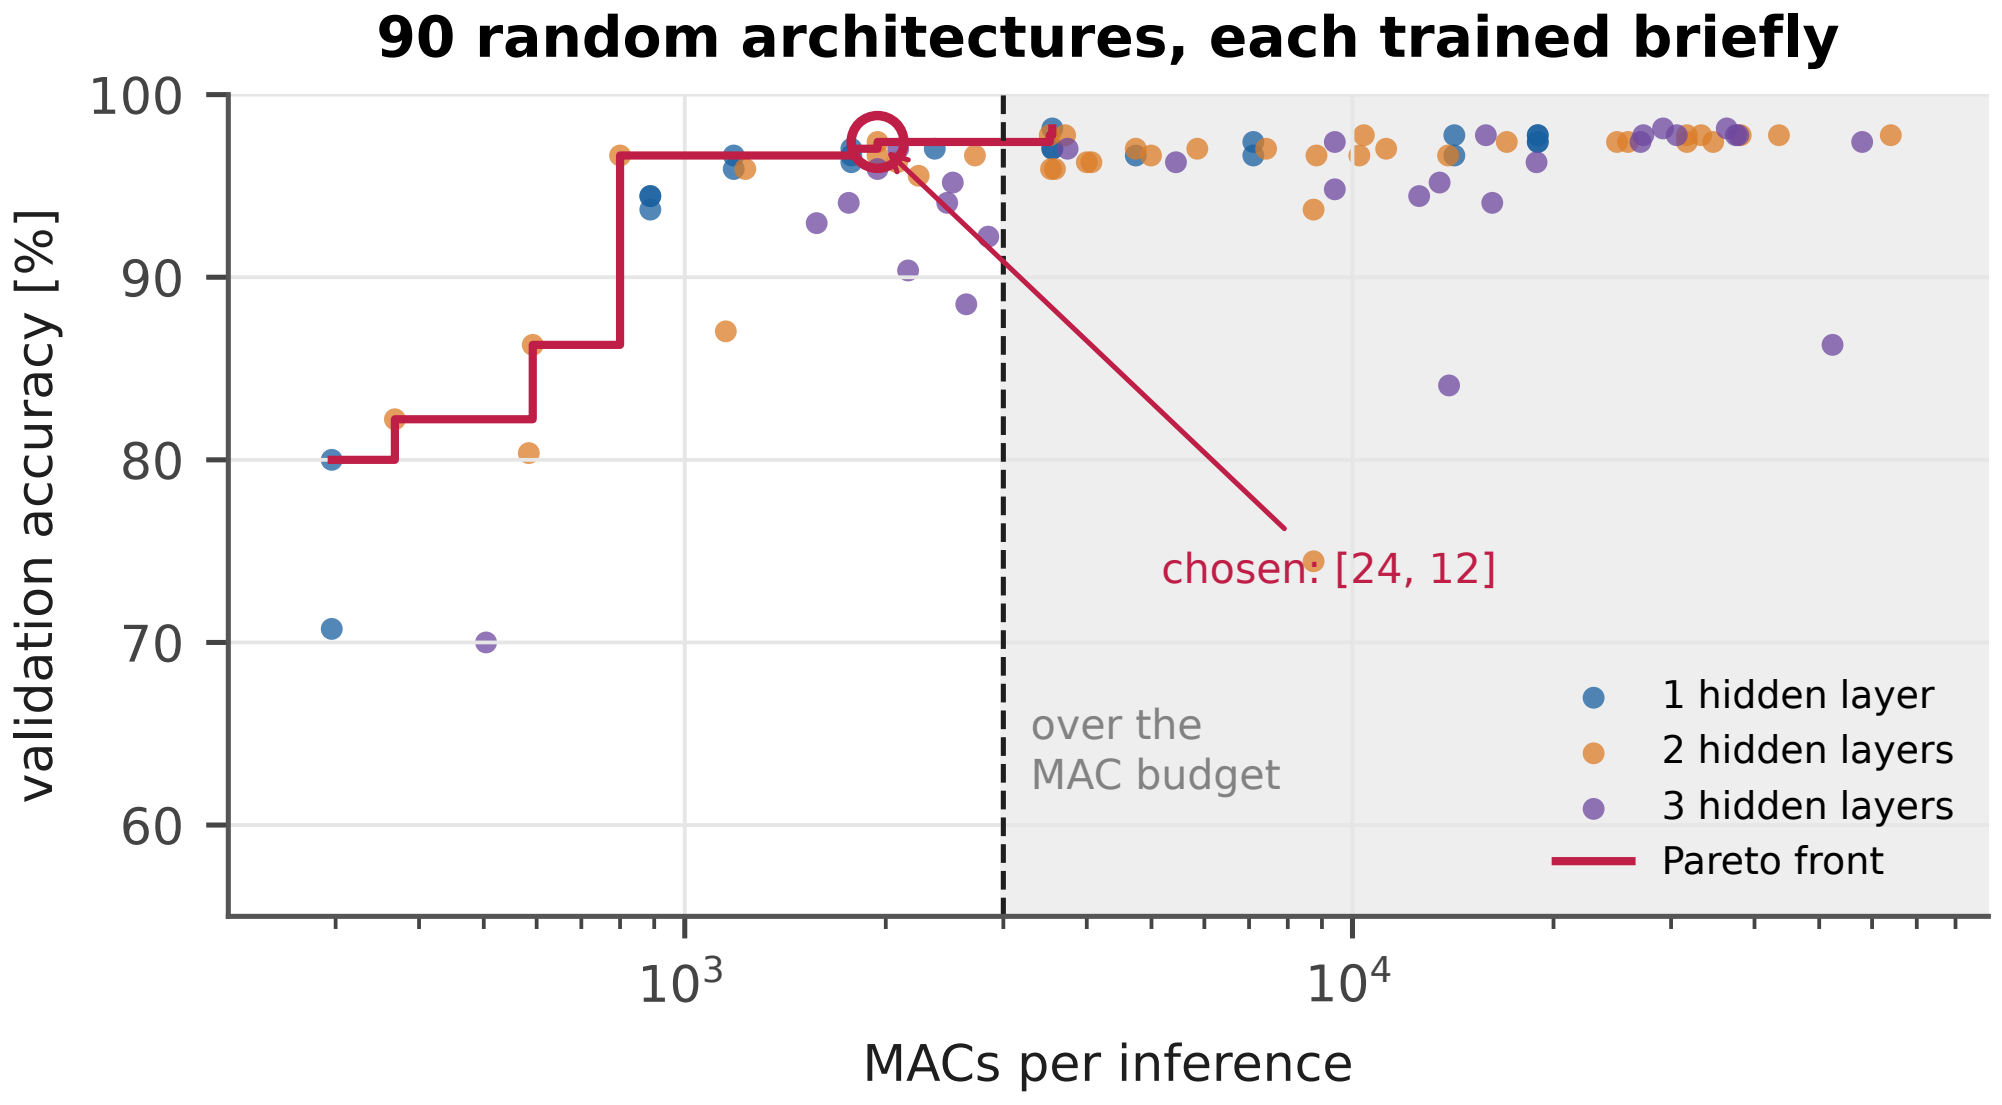

90 random architectures, each trained briefly
validation accuracy [%]
100
90
80
70
60
10^3
10^4
MACs per inference
chosen: [24, 12]
over the
MAC budget
1 hidden layer
2 hidden layers
3 hidden layers
Pareto front