



430 To test this, we evaluated three types of masks: preserved masks, globally permuted masks, and  
431 locally permuted masks. In the preserved mask case, the same mask found by the winning ticket is  
432 used for the random ticket. In the locally permuted case, the mask is permuted within each layer,  
433 such that the exact structure of the mask is broken, but the layerwise statistics remain intact. Finally,  
434 in the globally permuted case, the mask is permuted across all layers, such that no information should  
435 be passed between the winning ticket and the bad ticket.

436 Consistent with the lottery ticket hypothesis, we found that the winning ticket outperformed all  
437 random tickets (Figure A1). Interestingly, we found that while locally permuting masks damaged  
438 performance somewhat (blue vs. green), globally permuting the mask results in dramatically worse  
439 performance (blue/green vs. yellow), suggesting that the layerwise statistics derived from training the  
440 over-parameterized model are very informative. As this information would not be available without  
441 going through the process of generating a winning ticket, we consider the purely random mask to be  
442 the most relevant comparison to training an equivalently parameterized model from scratch.