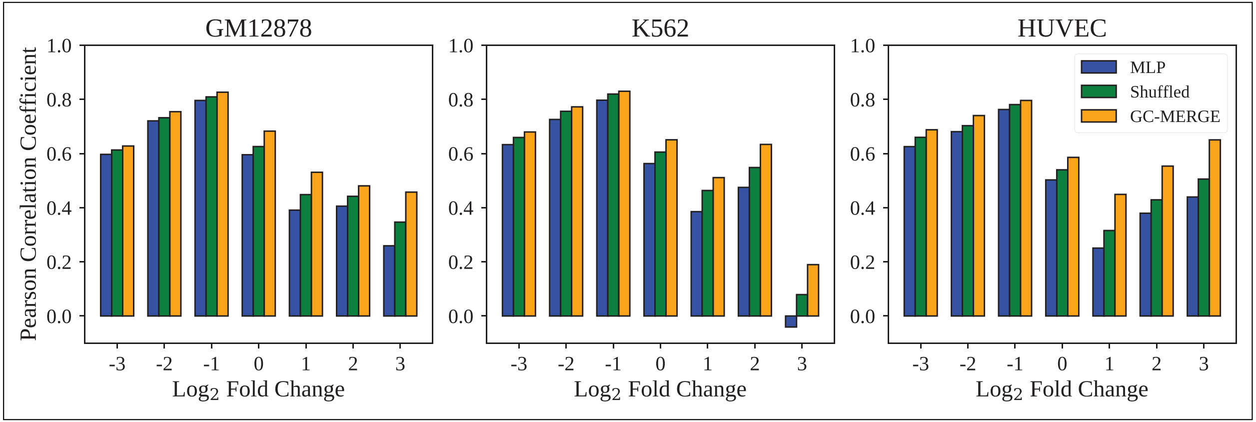

Figure S4: **Model performance broken down by level of differential expression.** For each cell line (GM12878, K562, and HUVEC), genes are stratified into subgroups binned by their log-2 fold change : $(-\infty, -3)$, $[-3, -2)$, $[-2, -1)$, $[-1, 0)$, $[0, 1)$, $[1, 2)$, $[2, 3)$, $[3, \infty)$. GC-MERGE consistently outperforms both the MLP and shuffled baselines, but the performance improvement is greater for genes with higher differential expression.