






S5 length generalization · three scoring rules

RLT-1 4+4 RLT-1 6+2 RLT-1 8+0
 RLT-1 5+3 RLT-1 7+1 Transformer 8



Seeds 42, 43, 44 · mean ± sample SD (n=3) · identical 1,024 sequences per length
 Panel-specific vertical scales · best ID-validation-loss checkpoint per run