



Task-free compression ladder: 24L → 8L → 6L → 4L



Sizes are unsharded safetensors bytes converted to MiB. KD losses are training objectives, not downstream scores.