

A212 · Only Two of Ten Nucleic Acid Foundation Models Encode Base-Pairing Partners

Presenter: Elliot Tower — University of Edinburgh

Authors: Elliot Tower

Session: Day 2 · Fri Oct 2, 2026 · 2:15–4:15 PM

Claims that RNA language models encode secondary structure rest on evaluations that do not control for composition. Stems are GC-rich and loops are AU-rich, so any composition-sensitive embedding inherits that enrichment as apparent structure awareness.

We apply a three-rung ladder to ten nucleic acid foundation models, five RNA-pretrained and five DNA-pretrained, across 47 Rfam families. Rung 1 asks whether mutating a base to its Watson-Crick complement perturbs stems more than loops; nine of ten models do, and a nucleotide-stratified null absorbs most of it. Rung 2 tightens the null to preserve dinucleotide composition, leaving ERNIE-RNA with 30 families and RiNALMo with 15, every other model below nine. Rung 3 asks whether a mutation perturbs its own base-pairing partner more than that partner's stem neighbors, against a null permuting partner assignments within each stem.

Two models resolve which position pairs with which. RiNALMo reaches 0.872 per-pair precision and ERNIE-RNA 0.888, where that null puts chance at 0.102 and 0.124. Choosing the readout layer on held-out families leaves these at 0.882 and 0.885 and places the other eight at or below their own chance rates. Pretraining domain does not order those eight: Evo at 7B parameters sits with the DNA-pretrained models.

A randomly initialized ERNIE-RNA clears every rung, reaching 0.490 against chance 0.241. Its architecture holds a hardcoded Watson-Crick table in a buffer that weight randomization does not reach. Ablating it drops the untrained model to its own chance rate while the trained model retains most of its excess.