BraintankAI

DEEP READ · GENOMIC MODEL BENCHMARKS · 7 MIN

Forty DNA models on one benchmark show why overall genomic AI rankings mislead

A single shared test of 40 genomic foundation models on 100 DNA tasks finds that overall leaderboards hide large differences between task types, and size explains only part of it.

13 November 2026 · arXiv preprint, not yet peer reviewed

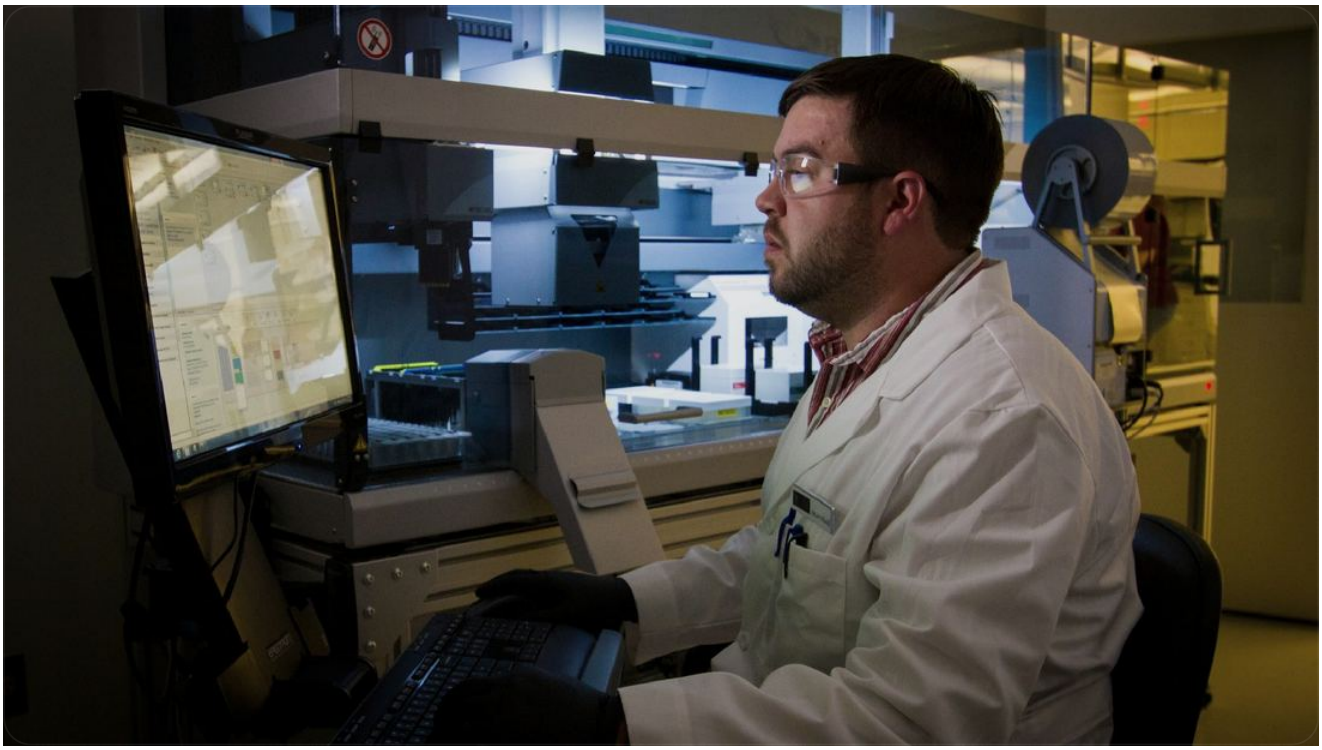


Photo: National Cancer Institute

Over the past few years, dozens of genomic foundation models have appeared. These are large neural networks pretrained on DNA sequences so they can be reused for many downstream prediction tasks. Each tends to arrive with claims of superior performance, but those claims are hard to check against one another. In a 2026 preprint, which has not yet been peer reviewed, Daria Ledneva, Mikhail Nuridinov and Denis Kuznetsov describe GENEB. It is a benchmark that runs 40 such models through 100 DNA classification tasks in 13 biological categories, all under one evaluation protocol. The authors state that none of the 40 models was developed by them or their funders.

The problem they set out to address is fragmentation. Genomic models are usually tested on different benchmarks, with incompatible data preparation and evaluation methods, and are often declared the best only in narrow settings. The authors illustrate this with a map of which models

cite which others as baselines. The map is sparse and disconnected. In their example, three prominent models, DNA-GPT, GenomeOcean and Evo, cannot currently be compared on principled terms. GENEB's main finding is that a single overall ranking conceals large differences from one category of task to another. Model size matters, but it does not fully account for those differences.

How it works

GENEB treats every model as a frozen feature extractor. Each DNA sequence goes into the model, which produces an embedding: an internal numerical summary of that sequence. The model itself receives no further training. The embeddings for each position along a sequence are averaged into one vector, a step called mean pooling. A simple classifier, logistic regression, is then trained on top of these vectors to make the task's prediction. This technique, known as linear probing, measures how much useful information the embeddings already hold. Because fine-tuning is left out, the result reflects the quality of each model's representation rather than how well it can be retrained.

Every task is run three ways. The probe is trained on one labelled example, on ten, or on the full labelled dataset, and each result is averaged over five fixed random seeds. The score is the Matthews Correlation Coefficient (MCC). It runs from −1 to 1, with 0 meaning chance, and unlike plain accuracy it is not inflated when one class dominates the data. Tasks with more than 100,000 sequences are cut down to 100,000. The authors justify this with an analysis of GenomeOcean embeddings, which showed scores levelling off beyond that size.

Scores are first averaged within each of the 13 categories, and those category scores are then averaged with equal weight. This "macro-average" is the headline metric. A "micro-average", which weights all 100 tasks equally and so favours categories with many tasks, is reported alongside it. To pull apart design choices, the authors also build controlled pairs: two models that differ mainly in one factor and are matched on the others. The factors are architecture, tokenisation (how DNA letters are split into units the model reads) and pretraining data.

What the evidence shows

There is no single baseline; the models are ranked against one another. Size clearly helps. The authors measure the Spearman correlation, a rank-based test of whether one quantity rises with another, between parameter count and overall score. It comes out at 0.573. Excluding Evo-1-131k, a model trained only on prokaryotic (bacterial-type) genomes, raises it to 0.694. Scaling is statistically significant in 11 of the 13 categories. Yet the advantage is modest and inconsistent. Models above a billion parameters beat those below 200 million by 0.075 MCC on average. The authors then look at 36 models, leaving out those pretrained only on prokaryotic, microbial or plant DNA. Among these, they count 25 cases where a model at least five times smaller beats a larger one. MutBERT, at 86 million parameters, outscores the 537-million-parameter eccDNAMamba by 0.110. Rankings barely change between macro- and micro-averaging, which correlate at 0.988.

The controlled pairs point to some design effects. Transformers beat eccDNAMamba, the one state-space model available for matched comparison. A state-space model is an alternative architecture that handles long sequences efficiently using a recurrent-style internal state. GenomeOcean-500M scores 0.550 against eccDNAMamba's 0.419. Transformer encoders, which read a whole sequence at once, beat decoders, which read left to right and predict the next unit, in all six matched pairs. The gaps are especially large on tasks involving other species: on virus and phage tasks, GenomeOcean-500M leads eccDNAMamba by 0.355. Chromatin accessibility, meaning whether a stretch of DNA is physically open to the cell's machinery, is the exception. There, eccDNAMamba leads by 0.124. No tokenisation scheme wins overall; most matched differences are within a few hundredths. Pretraining on many species beats human-only pretraining by only 0.018 overall, with most of the gain in a handful of categories. Models trained mainly on microbial DNA transfer poorly to the largely eukaryotic tasks, trailing by 0.084.

The starkest results concern scarce data. Across all 40 models, the average score falls from 0.488 with full data to 0.253 with ten examples and 0.106 with one. Three categories are close to random at one example: virus and phage, DNA methylation (chemical tags on DNA bases) and long non-coding RNA. Weaker models lose less in absolute terms, which the authors attribute to their low ceilings rather than genuine robustness. With little data, the best model changes in 8 of the 13 categories. Even with full data, 28 of the 100 tasks average below 0.35. Category leaders also differ from the overall leaders. Enformer ranks only moderately in aggregate but comes first in transcription-factor binding, enhancers and regulatory tasks. MutBERT is the strongest model at or below 100 million parameters, with an overall score of 0.529.

What it does not show

The authors list several limits. Very long-range tasks, over 10,000 DNA letters, are underrepresented. As a result, long-context models such as HyenaDNA-Large-1M and Evo-1-131k are not tested where their design should help most. Twelve of the 13 categories are eukaryotic, so the aggregate rankings systematically penalise prokaryotic and microbial models such as Evo-1-131k and DNABERT-S. Some tasks, especially DNA methylation and plant long non-coding RNA, may have noisy or weakly defined labels. Not every genomic model could be included, because of unavailable weights, incompatible pipelines or compute limits. Histone-modification and promoter tasks make up more than half the suite, which skews the micro-average toward them.

The protocol itself sets boundaries. Linear probing on frozen embeddings may underestimate what fine-tuning achieves. The authors check that rankings hold with a slightly more flexible classifier, a small neural network, but not against full fine-tuning. Mean pooling is used throughout. Its interaction with tokenisation is not separated out, even though letter-by-letter and fixed-length-chunk tokenisers produce much longer sequences than learned-chunk schemes.

Other gaps follow from the design. The conclusions about design factors rest on small numbers of matched pairs. The verdict on Transformers versus state-space models depends on a single model. The finding favouring curated eukaryotic-gene pretraining rests on one pair, which the authors flag themselves. The proposed link between chromatin accessibility and state-space designs is

tentative; the authors say only that the category may benefit from them. Many few-shot scores sit near the floor, and the authors describe one category's few-shot reranking as fragile to protocol choices. The subsampling threshold was set using one model's embeddings.

Where it fits

The authors compare GENEB to MTEB, a benchmark in natural language processing that brings many tasks together in one evaluation framework rather than a single-task leaderboard. Their sharper argument is against aggregate leaderboards in genomics. They contend that aggregate rankings in the low-data setting mix up how easy a category is with how good a model is. The category-by-category results are the evidence for that view. A model's place in the overall table says little about whether it leads on regulatory elements, methylation or viral DNA.

The broadest finding is descriptive: 40 models, 100 tasks, one protocol and five seeds. It is the most solid part of the work. Claims about which architecture, tokeniser or training corpus is better carry much less weight, because they rest on handfuls of matched pairs. The scaling story is best read as size mattering without being sufficient. The authors' own figures show a significant correlation overall and in most categories, alongside many cases of small models beating large ones.

As the opening story in this series on foundation models that read biology, GENEB serves as a map. It does not settle which genomic model is best. It shows why that question has rarely had a clean answer: rankings depend on the task category, on how much labelled data is available and on how evaluation is set up.

PAPER	*GENEB: Why Genomic Models Are Hard to Compare*
AUTHORS	Daria Ledneva, Mikhail Nuridinov, Denis Kuznetsov
VENUE	arXiv preprint — not yet peer reviewed · 3 June 2026
LINK	arxiv.org/abs/2606.04525
INDEX	stacks.braintank.ai
CODE	github.com/darlednik/GENEB