AlphaGenome 2 explained: predicting how DNA variants act in specific human cell types

24 November 2026 · 9 min read · blogs.iictp.org/alphagenome-2-cell-type-variant-effect-prediction-explained

Short answer: AlphaGenome 2 is a Google DeepMind model that predicts how DNA sequences and their variants affect gene regulation in specific human cell types, announced on 11 November 2026. DeepMind says it identifies nearly twice as many variants that change gene expression in cells as its previous state-of-the-art model. It matters technically because it moves from tissue averages to single-cell resolution, and from a fixed set of outputs to a model conditioned on a chosen biological state.

TL;DR The short version

- AlphaGenome 2, from Google DeepMind, predicts how DNA variants affect gene regulation in specific human cell types.
- DeepMind says it finds nearly twice as many variants that change gene expression as its previous best model.
- Predictions rank candidates but do not prove causation, so lab validation is still required.

Read online, listen to it, or change the text size: https://blogs.iictp.org/alphagenome-2-cell-type-variant-effect-prediction-explained

What is AlphaGenome 2?

According to DeepMind, AlphaGenome 2 is the next-generation version of AlphaGenome, which predicted molecular effects of DNA variants across broad tissues, and AlphaGenome Atlas, a platform with predictions for all 9 billion possible single-letter changes in the human genome. The new model zooms in to individual cell types and cell states. DeepMind says it is available for non-commercial research through the AlphaGenome API and in private preview for commercial research through Google Cloud.

How it works

- **A conditional model.** The first model read a one-million-letter DNA sequence and produced many tissue and cell-line outputs in parallel. AlphaGenome 2 instead takes the sequence plus one target condition, a tissue, cell type or state, as input at every stage, and predicts for that condition.
- **Cell encoder.** The condition can come from a reference catalogue of cell types or be generated from a scientist's own single-cell experiment through a cell encoder, so the model adapts to new data without retraining, according to DeepMind.

- **Cross-modality imputation.** Because the cell's state is separated from its outputs, a single measurement such as single-cell gene activity can be used to predict an unmeasured property such as chromatin accessibility.
- **Outputs.** It predicts gene expression, histone modifications and chromatin accessibility, and now DNA methylation, a widely requested readout.

Results reported by the source

- It was trained on bulk and single-cell data from human and mouse samples, including more than 100 million individual cell profiles across 68 tissues (the post also cites over 107 million profiles).
- On the TenK10K single-cell map of 28 primary immune cell types, DeepMind reports that AlphaGenome 2 identified causal variants at a rate of 30% versus 16%, and directional gene effects at 43% versus 23%, compared with the original AlphaGenome fine-tuned on single-cell data and other specialised tools.
- As a worked example, it reproduced a published finding in which a variant in a distant regulatory region drives inflammatory bowel disease by overactivating ETS2 in gut macrophages, singling out the causal variant and showing selective boosting of ETS2 in those cells.

Limitations and open questions

- These are DeepMind's reported results; the post links a paper, and peer-review status should be checked there.
- Predictions prioritise candidates; they are not evidence of causation. Lab validation is still required.
- Performance on cell states and diseases absent from training is described as a capability, but the post gives no tail-risk or failure analysis in the text we read.
- Access is limited to non-commercial research and a private commercial preview.

How to apply it in a research workflow

1. Start with the variants from a genome-wide association study or a rare-disease cohort and the cell types most relevant to the disease.
2. Run the variants through the model for those cell types and rank by predicted effect on nearby genes.
3. Compare against existing single-cell expression and accessibility data for the same cell types.
4. Take the shortlist to experimental validation, and record which predictions were wrong.

What it means for clinicians and biotech teams

Variant interpretation is a bottleneck in genetics. A cell-type-aware predictor does not change what can be done in a clinic today, but it shapes which hypotheses research teams test first. Clinicians should read it as a research tool until validated prospective use is shown. For another DeepMind and Google health model, see TB screening with foundation models.

Frequently asked questions

- **What is AlphaGenome 2?** A Google DeepMind model that predicts how DNA variants affect gene regulation in specific human cell types and states.
- **What is new compared with AlphaGenome?** It is conditional on a chosen cell type or state, trained on single-cell data, and adds DNA methylation as an output.

- **Is it available?** DeepMind says non-commercial research through the AlphaGenome API, and private preview for commercial research via Google Cloud.
- **Is it a diagnostic tool?** The announcement presents it as a research tool for prioritising variants.

Related reading

- Blind-sweep ultrasound AI and gestational age
- Smartwatch hypertension detection
- What is LLM hallucination?
- AI Practitioner

Sources and currency

Sources were read on 24 November 2026. Products, rules and standards change, so check the linked originals before you rely on any detail.

> Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Adam Gayoso, Natasha Latysheva, Pushmeet Kohli and Žiga Avsec, Google DeepMind. AlphaGenome 2: Deciphering how cells read DNA to better understand disease. 11 November 2026. deepmind.google/blog/alphagenome-2-deciphering-how-cells-read-dna-to-better-understand-disease.

Published by IICTP, an independent technical credential body, on 24 November 2026. The current version of this post is at https://blogs.iictp.org/alphagenome-2-cell-type-variant-effect-prediction-explained. More guides: blogs.iictp.org.