

News and Perspectives

AI Models Could Improve Diagnosis and Care for Rare Diseases

Simon Spichak, JMIR Correspondent

Abstract

Rare diseases are often critically underfunded, leaving many patients without timely diagnosis and treatment. In this *News and Perspectives* article, JMIR Correspondent Simon Spichak, who was a 2025 recipient of the National Press Foundation Rare Disease Reporting Fellowship, reports on advances in AI modeling that may offer new promise for rare disease detection and care.

Key Takeaways:

- Hundreds of millions of people worldwide who live with rare diseases face diagnostic delays and a lack of treatment options.
- AI models could help identify putative disease-causing variants.
- Frontier and foundation models may help speed up diagnosis and drug repurposing.

As many as [446 million people worldwide](#) are living with a rare disease. Because each affects relatively few individuals, affected individuals and their families often face years of [diagnostic delays](#), and [fewer than 5%](#) of rare diseases have approved treatments.

Researchers hope that advances in AI could speed up diagnosis and drug development. Specialized models trained on genomic information are already helping spot suspected disease-causing mutations, though these models are only good at very specific tasks.

Meanwhile, frontier and large foundation models built on vast biological datasets could help spot rare cancers and predict new treatments, though aside from a few success stories, they haven't yet proven their mettle in the real world.

Spotting New Protein-Coding Mutations

When a doctor suspects the cause of an unusual developmental disorder is genetic, they send their patient's DNA off for sequencing. Amino acids, the building blocks of proteins, are encoded by a triplet of genetic bases. But because of the built-in redundancy of this system, single base changes don't always result in a change to the amino acid or impact protein function.

As a result, DNA sequencing may identify over 10,000 protein-coding variants, but "most of these are benign," Rose Orenbuch, PhD, a postdoctoral fellow at Harvard University says. The challenge is predicting which of these changes actually impact the function of the protein built from these amino acids.

Several years ago, scientists trained an [AI model called EVE](#) on evolutionary information to predict how changes

in the genetic code affect protein function. But EVE had its limitations: the algorithm couldn't rank the pathogenicity of the mutations. That is, it would not be able to provide clinicians with enough information to determine which of these mutations is driving the rare disorder.

Orenbuch bridged this gap by training a new AI model called [popEVE](#). By adding a large language protein model onto EVE and training it on population-level human genetic data—even without additional information about the health of the individuals included in the dataset—it could determine which genetic changes might contribute to disease.

"We train these models unsupervised," says Orenbuch, "and we hope to find different patterns within that without actually needing the phenotype information." The algorithm doesn't receive any information about whether a mutation is associated with the disease. After the training, the popEVE model identified 123 novel genetic variants that may cause rare severe developmental disorders.

Orenbuch said that more than 25 of the genes have been verified by other researchers and added to a database of known disease-causing genes. The next step is developing models to identify when there are multiple disease-contributing variants occurring across the same pathway, as multiple mutations can sometimes be necessary to trigger a disease.

"They're not going to be a panacea," Orenbuch says. "These models are good for the data that we have, but there's always going to be an issue where you don't have high enough quality data for the patients."

Are Bigger Models Better?

Scientists are also training larger models to accelerate diagnosis and drug repurposing.

Built on the backbone of DeepSeek-V3, researchers from Shanghai's Artificial Intelligence Laboratory developed an agentic system called [DeepRare](#) that could help doctors diagnose rare diseases with traceable reasoning. Evaluated across 9 different datasets, the model's top-ranked prediction was correct 64% of the time, outperforming doctors and DeepSeek-V3 on its own.

In a [recent study](#) published in *NEJM AI*, another group of researchers used the OpenAI o3 Deep Research reasoning model to analyze genomic data from 376 unsolved rare disease cases. The researchers asked the model to propose the most plausible molecular explanation and provide its reasoning. The model was used as an aid, rather than a diagnostic tool. Its output, after being adjudicated by human experts, prompted physicians down the path for further testing and led to the clinical confirmation of 18 of these cases.

Others are moving away from frontier large language models to ones built on petabytes of biological information, called foundation models. [Anto](#) cofounder Arvid E Gollwitzer, PhD, is developing the backbone for foundation model development by making these databases easily searchable and providing a framework for training these large foundation models.

"They understand a specific domain really well—could be pathology images, genomics, proteomics, or even multi-omic data—and then they can be used for different downstream applications," Gollwitzer says.

“

*[These foundation models]
can help by transferring
knowledge learned from
large, diverse datasets into
settings where data is
scarce.*

Jean-Philippe Vert, PhD

Jean-Philippe Vert, PhD, is the cofounder and CEO of Bioprimus, which is developing foundation models. Such models could help researchers understand biology on an individualized level, he notes. By querying large biological datasets and compressing the patterns into a foundation model, "[these foundation models] can help by transferring knowledge learned from large, diverse datasets into settings where data is scarce," like in rare diseases, says Vert.

For example, since they affect less than [15 out of every 100,000](#) individuals, pathologists often struggle to make

an accurate rare cancer diagnosis. Researchers developed a foundation model called [Transformer-based pathology Image and Text Alignment Network](#) (TITAN), which was trained on more than 330,000 images of biopsies on slides. TITAN developed reports that could aid in diagnosis. Its abilities generalize to dozens of rare cancers, substantially outperforming other diagnostic aids.

Bioprimus has also developed and released free, open-source foundation models for understanding histopathology images. Vert says the models have been downloaded more than 1.5 million times and could help predict disease evolution and treatment response for tumors.

Another foundation model, [TxGNN](#), helps predict promising drug-repurposing candidates. Querying 3 different rare diseases, the model provides a ranked list of potential drugs for repurposing alongside an explanation of its reasoning.

Bioprimus's foundation models are also involved in drug development. "We distribute our more powerful models' commercial licenses to several of the top 10 pharma and major diagnostics companies," says Vert, to help accelerate and improve biomarker research, diagnostic tools, care paths, and clinical trials.

Ongoing Improvements and Next Steps for Implementation

Meanwhile, Gollwitzer is building the tools to ensure foundation models continue to improve. One of the limitations of current models is the sparsity of high-quality data, says Gollwitzer. Some models might not "find the underlying signal in these data sets, but get distracted by bias."

Anto developed a sparsification method that could pull out high-quality data from petabytes of publicly available datasets, discard the noise, and as a result, build stronger models. "If we train up a foundation model in this way, we're around 20 times more data-efficient than existing frontier models," says Gollwitzer.

Still, other researchers have pointed to [key constraints](#) of foundation models. These models may struggle to aggregate heterogeneous experimental and biological data. Many biobanks also struggle with attaining diversity in their samples, meaning that the data these models train on may not always be representative of the population. And while some models have shown promise, their ability to identify biomarkers and new treatments will need to be validated in future studies and trials.

Nonetheless, if these advances are validated and their constraints overcome, they could prove a boon for rare diseases.

Keywords: rare diseases; artificial intelligence; machine learning; diagnosis, computer-assisted; genetic variation; foundation models; frontier models

Please cite as:

Spichak S

AI Models Could Improve Diagnosis and Care for Rare Diseases

J Med Internet Res 2026;28:e106582

URL: <https://www.jmir.org/2026/1/e106582>

doi: [10.2196/106582](https://doi.org/10.2196/106582)

© JMIR Publications. Originally published in the Journal of Medical Internet Research (<https://www.jmir.org>), 16.Jul.2026