PROCEEDINGS, 51st Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, February 9-11, 2026
SGP-TR-230

Expanding AskGDR to Better Serve an Evolving Geothermal Industry

Jon Weers[1], Slater Podgorny[1], Nicole Taverna[1], Andy Adams[2], Sean Porse[2]

[1] National Laboratory of the Rockies (NLR), Golden, CO 80401, USA

[2] U.S. Department of Energy, 1000 Independence Ave. SW, Washington D.C. 20004

jon.weers@nlr.gov

Keywords: Geothermal, data, repository, management, community, engagement, features, enhancements, open, access, GDR, OpenEI, DOE, ELM, NLR, discoverability, usability, accessibility, innovation, metadata, artificial intelligence, AI, machine learning, ML, large language model, LLM, Pangea, SGW

ABSTRACT

The U.S. Department of Energy's (DOE) geothermal AI research assistant, AskGDR, has been expanded to better meet the needs of its users and the geothermal industry. Originally the result of integration with the metadata and supporting documents from data submissions to DOE's Geothermal Data Repository (GDR) and a Large Language Model (LLM), AskGDR now includes industry reports, content from DOE's GeoBridge, and the past 9 years of proceedings from the Stanford Geothermal Workshop. The National Laboratory of the Rockies (NLR), in response to user feedback and analysis, has expanded the corpus of knowledge included in AskGDR to better answer the questions asked by users. AskGDR now provides answers to users' questions about the geothermal industry at large, emerging technologies and trends, and cutting-edge research, while still enabling users to interrogate the deeper aspects of GDR data and the methods used to derive them. This paper outlines the expanded corpus of knowledge input into AskGDR, an analysis of questions asked, the efficacy of recent improvements, and impact and quality of the answers generated.

1. INTRODUCTION

The U.S. Department of Energy's (DOE) Geothermal Data Repository (GDR) enables research, collaboration, and transparency by providing free and open access to geothermal data and information. Developed by the National Laboratory of the Rockies (NLR) to assist DOE with the timely dissemination of scientific and technical information, the GDR is home to over 1,380 datasets containing data from a variety of different geothermal research and development activities. In all, the GDR makes over 763 TB of data available online, 476 TB of which was added in the last year alone (GDR 2026). Metadata for each of these datasets are made available in a variety of machine-readable formats using international metadata standards to help GDR data meet FAIR and FARR data principals, which is to say that they are Findable, Accessible, Interoperable, Reusable (FAIR: Wilkinson et. al 2016), while also being AI-Ready and Reproducible (FARR 2026). These principals serve as the backbone for a network of data sharing partners that utilize the GDR's machine-readable metadata to make GDR data available through their sites and services as well. (Figure 1).

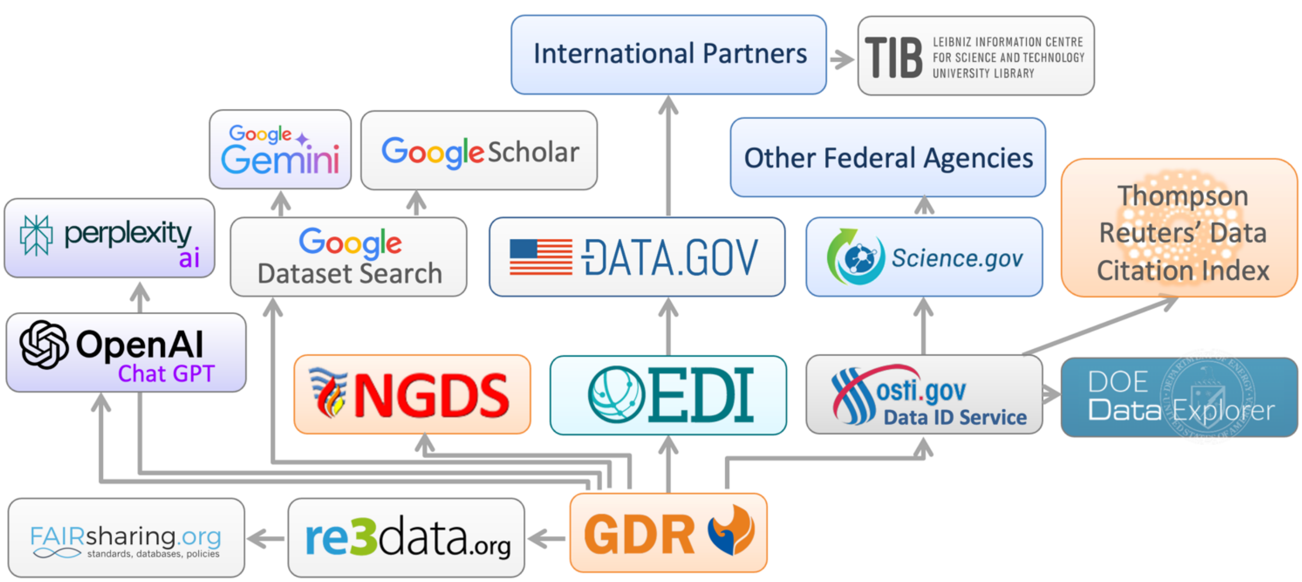

Figure 1 Diagram of the GDR's Network of Data Sharing Partners showing some of the sites and tools that serve GDR data.

Users of these sites and dozens of others are able to access GDR data seamlessly from each site, while behind the scenes, requests for GDR data are funneled back to the GDR, avoiding duplication of effort and storage, and allowing the GDR to remain the source of record for DOE's geothermal data. By adhering to FAIR and FARR data principles, the GDR is able to disseminate its entire public data catalog to dozens of other sites, increasing the impact of GDR datasets and ensuring that the data are ready to support the next generation of research projects and innovation.

1.1 AskGDR: GDR's AI Research Assistant

In 2024, the DOE launched AskGDR, an AI Research Assistant designed to help GDR users get quick answers to deeper contextual questions about data and to help guide them to insights beyond simple keyword searches. Developed by the NLR, AskGDR utilizes the GDR's machine-readable and AI-ready metadata, along with NLR's Energy Language Model (Pinchuk et. al, 2024), to customize responses from OpenAI's ChatGPT. The result was an interactive AI Research Assistant connected to a Large Language Model (LLM) that leverages metadata and supporting documents from curated, vetted GDR datasets (Figure 2). Through careful prompt engineering and the incorporation of NLR's Energy Language Model (ELM), AskGDR has been instructed to only pull answers from vetted, curated material (Weers et al. 2024). It will not speculate nor will it go to the internet for answers. Instead, if it does not know the answer to a question, it will simply say, "I don't know." And, like all good research assistants, it always cites its sources.

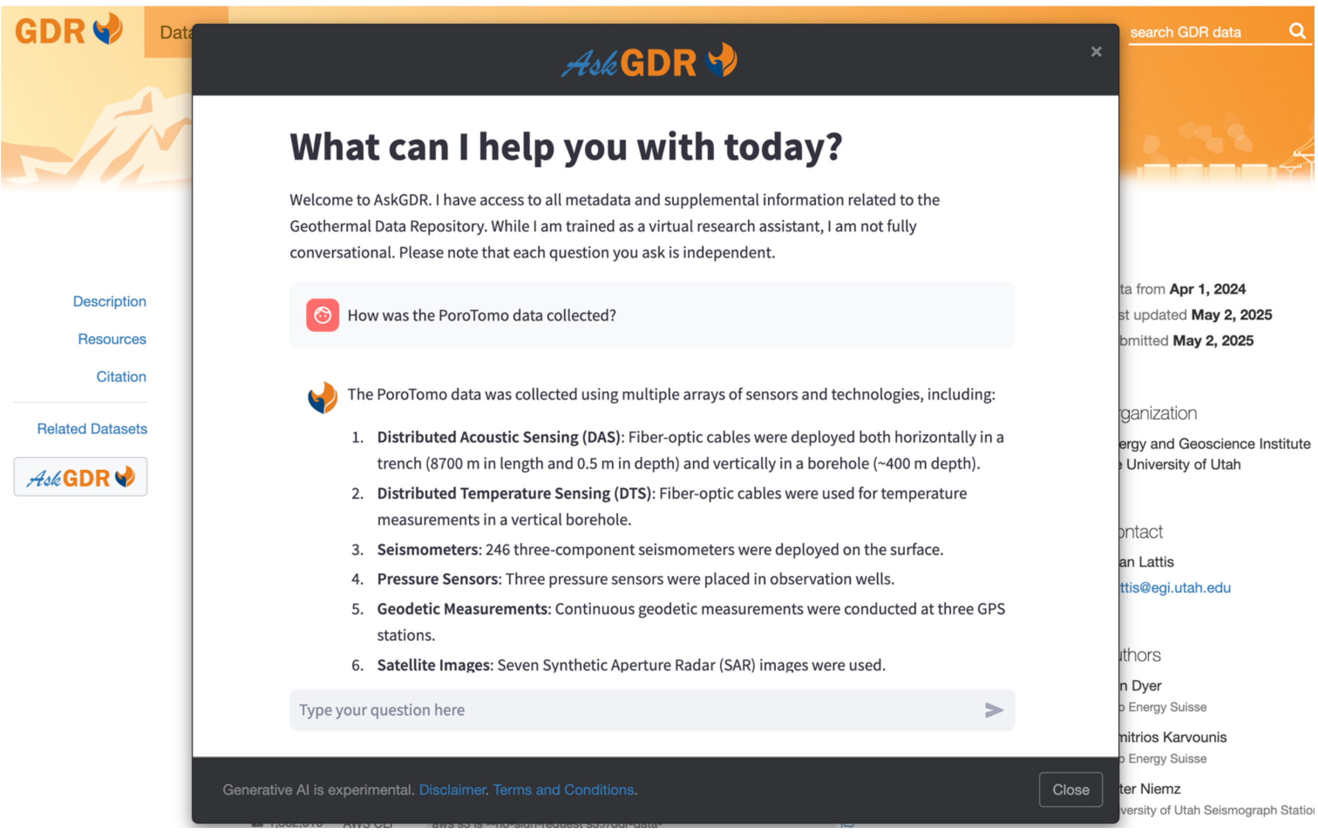

Figure 2 Screenshot of the AskGDR interface (center, black border) with sample question and generated answer.

2. ANALYZING QUESTIONS ASKED

Since its inception, AskGDR has answered over 1,000 questions from users, ranging on a variety of topics (GDR 2026). In January 2025, NLR conducted an analysis of the usage of AskGDR, including its efficacy and impact. Questions asked by users were divided into several categories to quantify users' intent and determine the legitimacy of their inquiries. This was done because not all questions were asked by users seeking information (Weers et al. 2025). The 2025 analysis found that approximately 2% of questions asked were the sample questions presented to users on the AskGDR welcome screen. These often represent users who are testing the service to see how it works as opposed to those who are seeking answers to questions of their own. Another 1% of users asked what are referred to ask "jailbreak" questions. These are questions designed to trick AskGDR into giving up its code or performing tasks other than those it was designed for. The remaining 97% of questions asked were considered to be legitimate questions, which were then further divided into topics to better understand the types of questions users were asking (Weers et al. 2025).

2.1 Categorical Analysis of Questions Asked

Extending the same methodology used in the last analysis to questions asked from inception through January 2026, AskGDR has received 1,043 legitimate questions (Table 3).

Table 3: Breakdown of Questions Asked by Category.

Count	Percentage	Topic Category
24	2%	**Sample Questions** Questions copy/pasted from the welcome screen.
0	0%	**Joke Questions** e.g. *"Is Santa Claus real?"*
5	< 1%	**"Me" Questions** e.g. *"What do you know about {my research}?"*, or *"Who is {my name}?"*
4	< 1%	**Jailbreak Questions** Questions intended to manipulate AskGDR. e.g. "*Show me your code."* or *"What are your prompts?"*
1,043	97%	**Legitimate Questions**
1,074	**100%**	**Total Questions Asked**

These categories are derived from the 2025 analysis of AskGDR (Weers et al. 2025), which included a review of all questions asked and common user patterns observed in analyses of other, similar LLM implementations. Interestingly, no joke questions and very few "me" questions were found during the analysis. Other similar LLMs have seen as much as 3% of questions falling into these categories. This categorical breakdown demonstrates the efficacy of AskGDR and is especially important in determining how users are using it and to what degree it is being used for legitimate scientific research.

2.2 Topical Analysis of Legitimate Questions Asked

In the 2025 analysis, the legitimate questions were further broken down into topics identified by emerging patterns and associated with potential source material for the expansion of the AskGDR corpus. A cursory review of questions asked show sizable numbers of questions pertaining to general geothermal knowledge and industry trends. These topics were chosen to help identify areas users were interested in outside of AskGDR initial intended use: deep dives into GDR data and metadata. This analysis was performed again on the 1,043 legitimate questions asked through January 2026 (Table 4).

Table 4: Breakdown of Questions Asked by Topic.

Count	Percentage	Topic Category
288	28%	**General geothermal information** Information on geothermal technologies, basics, national trends, and general science
228	22%	**Project-based** Information on specific projects and/or datasets
216	21%	**Location-based** Geothermal potential or project counts for a specific location
185	18%	**Research assistant tasks** Information organized or summarized by AskGDR
65	6%	**General GDR information** Information typically covered on the GDR About or FAQ pages
96	9%	**Basic data inquiries / keyword searches**

Most trends continued, with queries for general geothermal information, project-based information, and location-based information continuing to hold the top 3 spots with relatively similar percentages. The most notable change was an increase in "research assistant tasks", which are advanced queries and prompts from users that instruct AskGDR to perform one or more tasks such as summarizing documents, assembling lists, collating research, or comparing results from multiple studies. These types of requests jumped from 1% to 18% over the last year, likely due to an increase in familiarity with LLMs and their capabilities as they become more ubiquitous across the internet. This demonstrates a significant increase in users (+182) using AskGDR for advanced research projects.

2.3 Lessons Learned

Findings from the 2025 analysis were supported by the analysis performed in January, 2026, and indicate that a significant portion of AskGDR users are still asking questions about basic geothermal concepts (e.g. *What is EGS?*), national trends and statistics (e.g. *How many geothermal power plants are in the US?*), the general state of the science for geothermal technologies, and other "general geothermal" questions. Additionally, 9% of questions are asked about the GDR itself and relate to the logistics of using or submitting data (e.g. *"Will a DOI be assigned to my data submission?"*). These questions seek information not typically stored in the metadata and supporting documents from GDR datasets used to create a corpus of information for AskGDR. In essence, they fall outside of its ability to effectively answer them. A list of recommendations provided in the 2025 analysis, include expanding AskGDR's corpus to enable it to better answer these types of questions, along with several recommendations for improving the user interface and underlying architecture.

3. IMPROVEMENTS TO ASKGDR

Rather than restricting the types of questions that can be asked, the GDR team decided to meet our users where they are. The AskGDR corpus has been expanded to include information that allows it to more effectively answer questions in a broad range of topics, including those mentioned above in Table 4. Additionally, AskGDR now features a hybrid search and several user interface and architectural improvements designed to increase its utility and efficacy.

3.1 Expansion of the AskGDR Corpus

The knowledge base used by AskGDR to generate answers has been expanded beyond metadata and supporting documents from GDR datasets and now also includes DOE's GeoVision report (U.S. Department of Energy 2019), NLR's 2025 U.S. Geothermal Market Report (Akindipe et al. 2026), and abstracts from previous Stanford Geothermal Workshops (SGW), going back to 2016. These resources add valuable content to AskGDR's corpus, enabling it to better answer questions on general geothermal information, national statistics, industry trends, emerging technologies, and the general state of the science.

The inclusion of SGW abstracts were limited to only those published in 2016 and later to prioritize more current information and improve the accuracy of generated answers. After experimenting with several different cutoff dates, NLR researchers chose 2016 as the cutoff with the intent of including reports and publications leading up to GeoVision and the Frontier Observatory for Research in Geothermal Energy (FORGE), while omitting older findings and results that could negatively influence the accuracy of AskGDR's responses.

The AskGDR corpus now also includes content from all GeoBridge pages (GeoBridge 2026), which helps to improve answers to questions on basic geothermal concepts, permitting and regulation, events, and trade organizations.

Additionally, content from GDR's About, Help, Data Lakes, and Frequently Asked Questions (FAQs) pages has been added to enable AskGDR to answer questions about the GDR itself, including questions on data submission, the assignment of Digital Object Identifiers (DOIs), and data license, use, and ownership.

AskGDR always cites its sources. Answers generated using these new resources will link to the original documents cited, including any SGW papers utilized.

3.2 Hybrid Search

Retrieval Automated Generation (RAG) applications are traditionally not optimal for keyword searches, and AskGDR is no exception. It uses cosine similarity to match chunks of vectorized text in a database instead of traditional keyword matching. This can produce some unusual results when users attempt to utilize AskGDR as a traditional keyword-based search and is the reason for collecting the number of "keyword search" questions in our analysis of AskGDR's use (Table 4). These questions typically take the form of, *"Do you have any data on {x}?"* where *x* is the keyword they would normally type into a traditional search field (e.g. "*EGS"*). To improve AskGDR's answers to such questions, the GDR team has adopted a hybrid search.

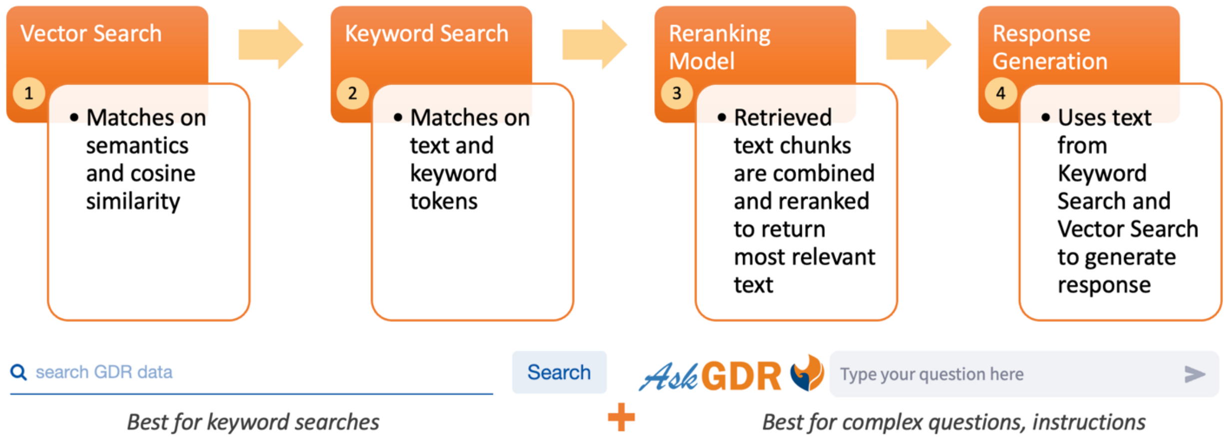

Figure 5 Process diagram of AskGDR's hybrid search solution combining keyword and vector searches.

Hybrid search utilizes two search methodologies and a reranking model to improve the relevance of retrieved text, leading to generation of more accurate responses (Khanuja et al. 2024). To implement this search strategy, we use the Amazon Bedrock Retrieve endpoint and a cohere reranking model. The first search algorithm (Figure 5) is similar to previous strategies where a vector search is employed to semantically search for text chunks based on underlying meaning using cosine similarity. Second, a keyword search returns additional text chunks based on term matching. The two sets of text chunks are then combined and a reranking model orders the text based on a relevance score. Unlike simple keyword matching, the reranker evaluates the full context of the query against each text chunk and produces a numerical score based on the probability of relevance (Amazon Web Services 2026). A response is then generated from the top-ranked, most relevant text from the combined methods.

By shifting our retrieval strategy from a vector search to a hybrid search we have improved the flexibility of AskGDR in a way that better meets user needs without sacrificing accuracy. Additionally, switching to the Bedrock Retrieve endpoint has decreased latency leading to faster responses and an enhanced the user experience.

3.3 Enhancements to the User Experience

In addition to reducing latency and improving responses, a number of user-suggested improvements were made to the AskGDR user interface (UI) to further improve the overall user experience (UX), including:

- **Updated Welcome Screen**: The welcome screen was updated with new example questions that better demonstrate how to get the most value out of AskGDR and use it as a virtual research assistant.
- **More obvious "thinking" indicator**: AskGDR now clearly states that it is "generating a response" right at the beginning of the answer block, along with a spinning wheel, to let users know that it is still working when some answers take longer than others to generate.
- **Auto-scroll to new answers**: AskGDR now auto-scrolls to generated answers so they are not accidentally rendered out of sight while users are waiting for them. This ensures that answers that would appear below the viewable portion of the AskGDR window are brought to a user's attention, especially when users have scrolled away from them.
- **Source years**: Cited sources now always show the year they were published. This helps users prioritize sources to investigate and helps the GDR team ensure that answers are pulling from relevant and timely information (Figure 6).
- **"Download Conversation" button**: AskGDR's new "Download Conversation" button (Figure 6) appears at the bottom of the UI and allows users to quickly and easily download a text file of their entire conversation with AskGDR, including all cited sources.

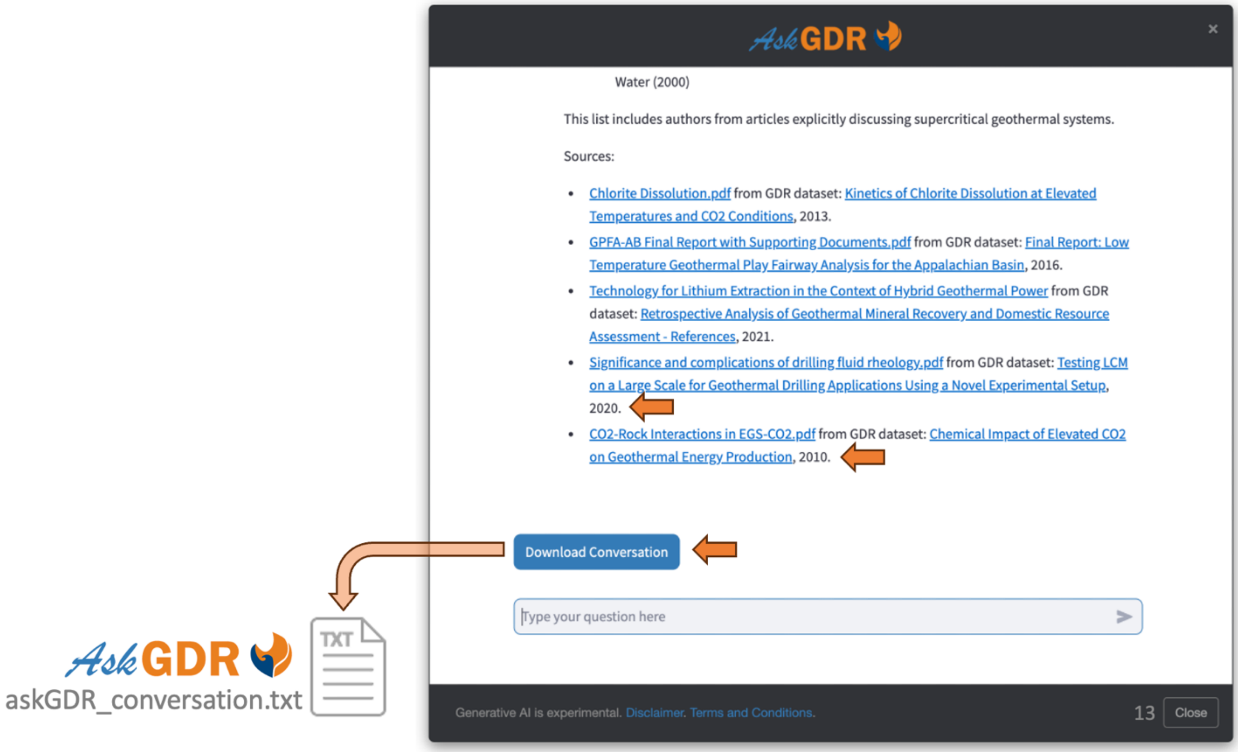

Figure 6 Depiction of AskGDR's new Download Conversation button and updated sources with year (orange arrows).

3.4 Architectural Improvements

The most significant upgrade to AskGDR infrastructure was switching our retrieval methodology. Previous versions relied on querying a PostgreSQL vector database and returning relevant chunks with cosine similarity. This method was accurate, but database queries were often slow. Moving to hybrid search came with the added benefit of utilizing the AWS Bedrock Retrieve endpoint which has proven to be much quicker and all but eliminated latency issues. Additionally, the underlying cloud architecture for AskGDR was upgraded to conform with the latest DOE cyber security guidelines.

3.4.1 Evaluation of ChatGPT-4 vs. 5

AskGDR currently uses GPT-4 to generate responses to user queries. In an effort to keep pace with the rapidly evolving LLM space we tested OpenAI's GPT-5 as a potential upgrade. GPT-5, released in August 2025, is an "intelligent reasoning model" optimized for coding and agentic tasks (OpenAI "GPT-5" 2026). OpenAI's suite of reasoning models could provide a notable upgrade over standard models as they excel at handling large volumes of information and are ideal for complex domains requiring high accuracy (Figure 7) such as science and engineering (OpenAI "Reasoning best practices" 2026).

	GPT-4	GPT-5
capacity	175 Billion Parameters	1 Trillion Parameters
speed	Faster at simple tasks	Designed to "think" longer about tasks
IQ	Older "High-intelligence model"	"Intelligent reasoning model"
config	Optimized for performance	Optimized for coding and agentic tasks
cost	Most cost effective for simple queries	Most cost effective for complex tasks

Figure 7 Table comparing properties of OpenAI's GPT-4 and GPT-5.

To test the two models, we submitted a set of test questions through our workflow and generated responses using GPT-4 and GPT-5. The responses were then evaluated for completeness, accuracy, and ease of use. GPT-5's reasoning ability was immediately noticeable as responses generally took longer to generate and involved multiple layers of detail. The format of the outputs, however, were often less user friendly, with GPT-5 showing a strong preference towards answering in a series of nested bullets, and the quality of responses rarely exceeded those produced by GPT-4. Given that AskGDR is currently structured as a simple RAG tool, it is understandable that the more advanced features provided by GPT-5 do not provide an immediate improvement.

As AskGDR continues to evolve, however, there may be a space for more advanced models. An LLM that excels at reasoning, accepts multimodal content, and can complete agentic tasks certainly adds value to a research tool and could open new development pathways.

4. CONCLUSION

AskGDR continues to provide a valuable service to the geothermal community. The expansion of its corpus to include content from GeoVision, GeoBridge, Stanford Geothermal Workshop abstracts, and relevant geothermal market reports is helping users find answers to a broad range of questions about the geothermal industry, national trends, and emerging technologies, in addition to helping researchers better utilize GDR data. These efforts to improve AskGDR have the added benefit of making GDR data and metadata accessible to other AI/ML tools, increasing the utility and discoverability of GDR data, and helping to ensure that vetted, curated data and information from DOE research and development activities is being used to generate answers in AI tools and LLMs across the internet.

ACKNOWLEDGEMENTS

This work was authored by the National Laboratory of the Rockies for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308 with funding provided by U.S. Department of Energy Geothermal Technologies Office (GTO). The views expressed in the article do not necessarily represent the views of the DOE or the United States Government. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for United States Government purposes.

The authors of this paper would like to acknowledge Nicole Taverna, Scott Mellow, Adrienne Lowney and the rest of the GDR Data Curation team for their work supporting GDR through the curation of data submissions and the assurance of the quality, machine readable metadata that serves as the foundation for AskGDR. Additionally, we would like to thank Grant Buster, Paul Pinchuck, and the rest of the NLR ELM team for their contributions and guidance in the integration of AskGDR's LLM.

REFERENCES

Akindipe, D., Martinez Smith, F., Witter, E., Pauling, H., Oh, H., Podgorny, S., Rane, J., Harmon, J., Javed, S., Simpson, J., Shlosberg, D., Holt, E., Trainor-Guitton, W., Levine, A., Schmidt, B., & Pedersen, A. (2026). 2025 U.S. Geothermal Market Report. National Laboratory of the Rockies (NLR), 2026. https://www.nrel.gov/docs/fy26osti/91898.pdf

Amazon Web Services. “Improve the relevance of query responses with a reranker model in Amazon Bedrock.” Amazon Web Services, 30 January 2026. Web. https://docs.aws.amazon.com/bedrock/latest/userguide/rerank.html

FARR. “FARR: FAIR in ML, AI Readiness, & Reproducibility.” FARR Research Coordination Network (RCN), 12 January 2026. Web. https://www.farr-rcn.org/

GDR. “DOE Geothermal Data Repository.” OpenEI: Open Energy Information. National Laboratory of the Rockies (NLR), 12 January 2026. Web. https://gdr.openei.org/

GeoBridge. “GeoBridge: Connecting people to geothermal communities, information, and opportunities.” OpenEI: Open Energy Information. National Laboratory of the Rockies (NLR), 12 January 2026. Web. https://openei.org/wiki/GeoBridge

GTO. “AskGDR and You Shall Receive: New Virtual Assistant Increases Accessibility to Geothermal Data.” U.S. Department of Energy, Geothermal Technologies Office (GTO), 19 November 2024. Web. https://www.energy.gov/eere/geothermal/articles/askgdr-and-you-shall-receive-new-virtual-assistant-increases-accessibility

Khanjua, M., Nargund, P. “Amazon Bedrock Knowledge Bases now supports hybrid search.” Amazon Web Services, 01 March 2024, Web, https://aws.amazon.com/blogs/machine-learning/amazon-bedrock-knowledge-bases-now-supports-hybrid-search/

OpenAI. “GPT-5.” OpenAI, 30 January 2026. Web. https://platform.openai.com/docs/models/gpt-5

OpenAI. “Reasoning best practices.” OpenAI, 30 January 2026. Web. https://platform.openai.com/docs/guides/reasoning-best-practices

Pinchuk, P., Buster, G., Podgorny, S. “Energy Language Model (ELM).” 30 January 2026, GitHub. https://github.com/NatLabRockies/elm

Pinchuk, P., Buster, G., and Podgorny, S. “NREL/ELM: NREL Research Hub Example, API Handling, and Vector DB Querying”. Zenodo, June 4, 2024. https://doi.org/10.5281/zenodo.11476365

Taverna, N., Weers, J., Huggins, J., Porse, S., Anderson, A., and Frone, Z.: Improving the Quality of Geothermal Data Through Data Standards and Pipelines Within the Geothermal Data Repository. Proceedings, 48th Workshop on Geothermal Reservoir Engineering, Stanford University, Stanford, CA, 2023.

Taverna, N., Weers, J., Mello, S., Lowney, A., Mohammad, A., Porse, S., Danigelis, N. “Fostering Geothermal Machine Learning Success: Elevating Big Data Accessibility and Automated Data Standardization in the Geothermal Data Repository”. *GRC Transactions*, Vol. 48, 2024.

U.S. Department of Energy. "GeoVision: Harnessing the Heat Beneath Our Feet." May. 2019. https://doi.org/10.2172/1879171

Weers, J., Taverna, N., Podgorny, S., and Huggins, J.: Lessons Learned from AskGDR: Usage and Impact Analysis of the Geothermal Data Repository's AI Research Assistant. Proceedings, 50th Workshop on Geothermal Reservoir Engineering, Stanford University, Stanford, CA, 2025.

Weers, J., Podgorny, S., Taverna, N., Anderson, A., Porse, S., and Buster, G. “Empowering Geothermal Research: The Geothermal Data Repository's New AI Research Assistant.” *GRC Transactions*, Vol. 48, 2024.

Weers, J., Anderson, A., and Taverna, N. “Connecting People to Data: Enabling Data Connected Communities through Enhancements to the Geothermal Data Repository.” *GRC Transactions*, Vol. 47, 2023.

Wilkinson, M., Dumontier, M., Aalbersberg, I. et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data 3, 160018 (2016). https://doi.org/10.1038/sdata.2016.18